Empirical Informational Distortion of Map Projections

A Cross-Regional Study on Real Settlement Distributions

A companion volume to Original Mathematical Perspectives on Cartography

Silvano Salvador

Contents

Abstract **iv**

Note on This Extended Printing **v**

1 Introduction: From a Fragile Asymptotic Claim to a Robust Empirical One **1**

1.1 Where the parent volume left off . . . 1

1.2 The analytic route, and why it is reported as a failure . . . 1

1.3 The empirical route, pursued here . . . 2

1.4 What this volume does, and does not, claim . . . 3

1.5 Summary . . . 3

2 Methodology: From Synthetic Grids to Real Settlement Data **4**

2.1 The GeoNames gazetteer . . . 4

2.2 Sampling scheme . . . 4

2.3 Reuse, unchanged, of the parent volume's IDI pipeline . . . 5

2.4 Projection centring for national-scale analysis . . . 6

2.5 Aggregation across countries within a region . . . 6

2.6 Summary . . . 7

3 Mediterranean Europe: Italy, Spain, Portugal, Croatia, Slovenia, France **8**

3.1 Regional character . . . 8

3.2 Regional summary . . . 8

3.3 Country-by-country detail . . . 8

3.4 Interpretation . . . 11

3.5 Summary . . . 12

4 The Maghreb: Algeria, Morocco, Tunisia, Libya, Mauritania **15**

4.1 Regional character . . . 15

4.2 Regional summary . . . 15

4.3 Country-by-country detail . . . 15

4.4 Interpretation . . . 17

4.5 Summary . . . 19

5 The Caucasus and Anatolia: Armenia, Azerbaijan, Georgia, Turkey **21**
5.1 Regional character . . . 21
5.2 Regional summary . . . 21
5.3 Country-by-country detail . . . 22
5.4 Interpretation . . . 23
5.5 Summary . . . 25

6 South Asia: India and Pakistan **26**
6.1 Regional character . . . 26
6.2 Regional summary . . . 26
6.3 Country-by-country detail . . . 27
6.4 Interpretation . . . 28
6.5 Summary . . . 29

7 Atlantic and Western Europe: Austria, Switzerland, the United Kingdom, Ireland **30**
7.1 Regional character . . . 30
7.2 Regional summary . . . 30
7.3 Country-by-country detail . . . 30
7.4 Interpretation . . . 31
7.5 Summary . . . 33

8 Synthesis: One Robust Regularity Across Twenty-One Countries **35**
8.1 The result, stated once, plainly . . . 35
8.2 Cross-regional rank stability . . . 36
8.3 Regional variation in the grand-mean effect size . . . 36
8.4 The trio's internal ordering, country by country . . . 37
8.5 What drives the effect, and what does not . . . 39
8.6 Relation to the failed scaling-law attempt . . . 39
8.7 Summary and closing remark . . . 39

9 Robustness Diagnostics **41**
9.1 Bin-count sensitivity . . . 41
9.2 Sample-size and sample-composition sensitivity . . . 43
9.3 Cross-regional rank concordance . . . 45
9.4 Country shape and the LCC/LAEA margin . . . 45
9.5 Gazetteer size and the trio-completion gap . . . 47
9.6 Illustrative distance distributions . . . 48

9.7 Summary 49

10 The Anisotropy Scaling-Law Attempt: A Documented Failure 52
10.1 The prediction 52
10.2 The test, and its result 52
10.3 Diagnosis 54
10.4 Why this is reported rather than discarded 54

11 Computational Implementation and Reproducibility 55
11.1 The empirical IDI pipeline 55
11.2 LaTeX table generation 56
11.3 The anisotropy scaling-law test 56
11.4 Robustness diagnostics 56
11.5 The trio-completion-gap correlation 57
11.6 Extended robustness diagnostics 57
11.7 Verification discipline 58

12 Five Additional Projections 59
12.1 Selection rationale 59
12.2 Formulas and verification 59

13 Extending the Synthesis: Seventeen Projections 61
13.1 The grand mean, recomputed 61
13.2 Regional tables 61
13.3 Cross-regional rank stability 62
13.4 Where the trio's exceptions and the quintet's exceptions coincide 63

14 Extended Robustness Diagnostics: The Seventeen-Projection Quintet 68
14.1 Bin-count sensitivity, seventeen projections 69
14.2 Sample-size and sample-composition sensitivity, seventeen projections 71
14.2.1 Locating the crossover 73
14.2.2 Narrowing the bracket, and a control pair that does not cross 74
14.3 Elongation and the Bonne/Polyconic margin 76
14.4 Gazetteer size and the quintet-completion gap 78
14.5 Summary 79

References 81

Abstract

The companion monograph *Original Mathematical Perspectives on Cartography* introduced the Informational Distortion Index (IDI) — a Jensen–Shannon divergence between the distribution of true geodesic distances and the distribution of projected planar distances over a sampled point set — and validated it exclusively on systematically constructed latitude/longitude grids. This volume asks the natural follow-up question: does the same diagnostic, applied not to synthetic grids but to *real settlement distributions*, tell a coherent and useful story? An initial attempt to answer a more ambitious question — whether a clean analytic scaling law governs the IDI as a sampled region shrinks to a point, tested numerically in Chapter 10 — did not survive numerical scrutiny: the histogram-based JSD estimator saturates in a combinatorial noise floor before any asymptotic regime is reached, and the predicted exponent and ratios are not recovered at any of the six angular scales tested, spanning a 32-fold range from 20° down to 0.625°. Rather than force that result into print, this volume redirects to a more modest and more robust question. Using real populated-place coordinates from the GeoNames gazetteer for twenty-one countries across five geographically and culturally diverse regions (Mediterranean Europe, the Maghreb, the Caucasus and Anatolia, South Asia, and Atlantic/Western Europe), the twelve projections already implemented in the parent volume are re-evaluated against genuine, non-uniform, population-concentrated point clouds. The result is a single, robust, falsifiable empirical regularity: at national and regional scale, projections centred and fitted to the region (Lambert Azimuthal Equal-Area, Albers Equal-Area Conic, Lambert Conformal Conic) systematically outperform projections optimized for whole-Earth compromise display (Robinson, Mollweide, Hammer, Winkel Tripel, Van der Grinten, Aitoff) by two to three orders of magnitude in IDI, with no exception across all twenty-one countries tested. All numerical results reported here are the direct, unedited output of the pipeline in Chapter 11, executed on the real GeoNames data supplied by the author; no figure is invented or retrofitted, and the one substantial theoretical attempt that failed is reported as a failure, not quietly discarded.

Note on This Extended Printing

The abstract above, and Chapters 3 to 11, are unchanged from the volume's original printing and report exactly what they always reported, on twelve projections. Chapters 12 and 13 are new: they introduce five further projections (Gall–Peters, Eckert IV, Eckert VI, Bonne, Polyconic) and rerun the entire twenty-one-country pipeline once, from scratch, scoring all seventeen projections together. The extension does not revise any number in the original chapters — the twelve-projection results stand exactly as published — and its own headline finding is a strengthening, not a qualification, of the original one: the three-member regionally-fitted trio turns out to be the leading edge of a five-member quintet, with no new exception introduced anywhere the original study did not already report one. Chapter 14 is newer still: it puts the seventeen-projection extension through the same style of robustness scrutiny Chapter 9 already applied to the twelve-projection original, reruns the two arbitrary-parameter sensitivity checks and the gazetteer-size null result against the quintet, and adds one further check — a country-shape correlation against the Bonne/Polyconic margin — that has no twelve-projection analogue because neither projection existed before Chapter 12. It finds the quintet's occupation of the leading ranks exactly as robust as the trio's was, but its internal ordering measurably less so, and reports both results, together with two still-open correlation questions, without smoothing either into a stronger claim than the data supports.

Chapter 1

Introduction: From a Fragile Asymptotic Claim to a Robust Empirical One

1.1 Where the parent volume left off

The closing chapter of *Original Mathematical Perspectives on Cartography* listed, under the heading "Directions not pursued here," several concrete extensions of the Informational Distortion Index (IDI) that the author had deliberately set aside: a rigorous treatment of Chebyshev optimality on general (non-linearized) regions via the Beltrami equation and quasiconformal maps; the interaction of the Hairy Ball obstruction with higher-order vector data such as vorticity; and, most relevant here, the observation that the IDI had so far been computed *only* on systematically constructed latitude/longitude grids, never on the real, non-uniform, population-concentrated point distributions that any actual map is made to serve. That final item is the starting point of this volume.

Two distinct routes suggested themselves for pursuing it further. The first, more ambitious route asked whether the local behaviour of the IDI, as the sampled region shrinks toward a single point, obeys a clean analytic scaling law derivable from the second-order Taylor expansion of the azimuthal projection family in geodesic normal coordinates — a genuine piece of new mathematics, in the spirit of Milnor's quantitative treatment of Chebyshev's theorem already used in Chapter 3 of the parent volume. The second, more modest route set the analytic question aside entirely and asked instead a purely descriptive question: computed on *real* settlement data rather than synthetic grids, does the projection ranking implied by the IDI tell a stable, non-arbitrary story?

1.2 The analytic route, and why it is reported as a failure

The analytic route was pursued first, and is documented in full in Chapter 10, because it is instructive precisely in how it fails. A second-order expansion of the sphere's metric in geodesic normal coordinates around a point p_0 (Bertrand–Diguet–Puiseux), compared against the corresponding second-order expansion of an azimuthal projection's radial scale function $k'(c) = 1 + ac^2 + O(c^4)$, yields a clean, closed-form invariant

$$\text{Aniso}(a) := (\text{radial coefficient}) - (\text{tangential coefficient}) = 4a - 1, \tag{1.1}$$

which vanishes exactly for the stereographic projection (the unique azimuthal projection that is conformal, hence isotropic to second order at its centre) and is nonzero for the gnomonic, Lambert azimuthal equal-area, and azimuthal equidistant projections. This calculation is correct and was verified line by line; the difficulty appears only in the next step, extrapolating from this pointwise metric statement to a prediction about the IDI, a JSD computed over an aggregate distribution of pairwise distances across a whole sampled disc. The extrapolation predicts $\text{IDI}(f, \Omega_\eta) \sim C_f\, \eta^4$ as the angular radius $\eta \to 0$, with $C_f \propto \text{Aniso}(a_f)^2$.

Methodological Note

This prediction was tested numerically *before* any chapter text was written around it, using an isotropic (sunflower-arrangement, not regular-grid) sampling scheme of 169 points inside a shrinking disc, exactly as detailed in Chapter 10. It failed cleanly and reproducibly: the observed log-log scaling exponent does not stabilize near 4, or indeed near any single value — it ranges from about 3.4 down to about 0.2 across the four projections and six tested scales, generally trending *downward* rather than converging as the disc shrinks, and collapsing to an undefined slope at the smallest scale once the JSD saturates at the fixed histogram-resolution floor. The predicted asymptotic ratio of IDI values across projections is likewise not recovered at any tested angular scale. The most likely diagnosis, argued in detail in Chapter 10, is that the relevant asymptotic expansion holds for the Fisher-Rao quadratic approximation applied pointwise, but not for the JSD of two *aggregate* distance distributions with slightly mismatched supports, where a boundary effect of order η^2 rather than η^4 dominates, compounded at the smallest scales by the same fixed-bin-count saturation already documented for national-scale data in Chapters 5 and 7. This has not been proved, only diagnosed as the most likely mechanism consistent with everything observed. The chapter is not written around this claim.

Reporting a negative result this way — with the calculation that holds kept separate from the extrapolation that does not, and the failure mode diagnosed rather than hidden — follows the same standard of transparency already established in the parent volume's Chapter 9, which reports both the Aitoff/Hammer ranking instability at the top of the global-grid ordering and the non-significant Kendall's τ correlation across sampling regions exactly where they occur rather than smoothing them into a cleaner-sounding story. A monograph that reports only confirmations is not more rigorous for it.

1.3 The empirical route, pursued here

Setting the fragile asymptotic claim aside, this volume pursues the second, more modest route to completion: apply the existing, already-validated IDI machinery — unchanged in its definition, unchanged in its twelve projection formulas — not to synthetic grids but to real populated-place coordinates drawn from the GeoNames gazetteer, across a geographically and culturally diverse set of twenty-one countries grouped into five regions:

- **Mediterranean Europe**: Italy, Spain, Portugal, Croatia, Slovenia, France.
- **Maghreb**: Algeria, Morocco, Tunisia, Libya, Mauritania.
- **Caucasus and Anatolia**: Armenia, Azerbaijan, Georgia, Turkey.
- **South Asia**: India, Pakistan.
- **Atlantic / Western Europe**: Austria, Switzerland, the United Kingdom, Ireland.

Methodological Note

The fifth group is included for an openly stated reason of data availability rather than thematic fit: the country roster assembled for this study (originally gathered for the author's independent toponymic-entropy work) includes several Atlantic and Central European countries — Austria, Switzerland, the United Kingdom, Ireland — that are neither

Mediterranean, Caucasian, South Asian, nor Maghrebi. Rather than force them into one of the four thematically motivated groups, where they do not belong geographically or climatically, they are kept as a fifth, explicitly labelled comparison group. This is exactly the kind of methodological honesty the parent volume insists on elsewhere: no dataset is reshaped to fit a narrative that does not describe it.

1.4 What this volume does, and does not, claim

This is a descriptive and comparative study, not a new theorem. It asks a single question separately in twenty-one real national settings: when a country's actual populated places — not an abstract lattice — are projected by each of twelve standard projections, which projections best preserve the true geodesic distance structure between those places, measured by the IDI? The answer, detailed region by region in Chapters 3 to 7 and synthesized in Chapter 8, is unambiguous at first place and holds without a single exception across all twenty-one countries: a projection fitted to the region beats every projection designed for whole-Earth compromise display, typically by one to two orders of magnitude, regardless of latitude, territorial shape, or settlement density. This is a modest claim by design. It requires no new mathematics, carries essentially no risk of being an artifact of small-sample noise (each country contributes $\binom{80}{2} = 3160$ pairwise comparisons, none of them synthetic), and is fully reproducible from the pipeline in Chapter 11 applied to the raw GeoNames extracts the author supplied. One honest qualification belongs here rather than being smoothed away in the regional chapters: *first* place goes to a regionally-fitted projection (Lambert Azimuthal Equal-Area, Albers, or Lambert Conformal Conic) in all twenty-one countries without exception, but the stronger claim that all three regionally-fitted projections always occupy the top three ranks does not hold universally — it fails, in a specific and reproducible way, in Spain, Algeria, and India, where the Sinusoidal or Winkel Tripel projection displaces the weakest member of the trio out of third place. Chapters 3, 4 and 6 report this exactly where it occurs, rather than rounding it off.

1.5 Summary

The parent volume's Chapter 11 identified real, non-grid data as unfinished business. This volume closes that item, not by proving the ambitious analytic conjecture attempted first (documented honestly as a failure in Chapter 10), but by carrying out the more modest empirical programme to completion, on real data, with a result that is small in scope but solid in every part that was actually tested.

Chapter 2

Methodology: From Synthetic Grids to Real Settlement Data

2.1 The GeoNames gazetteer

GeoNames is a large, freely available, community-maintained geographical database, distributed as per-country tab-separated text files. Each record includes, among other fields, an ASCII place name, latitude and longitude in decimal degrees, a one-letter feature class, a three- or four-character feature code, an ISO country code, first-order administrative subdivision, and a population figure (zero or blank for many minor features). Feature class `P` denotes populated places — cities, towns, villages, and administrative capitals — as opposed to hydrographic features (`H`), terrain features (`T`), undersea features (`U`), and the other seven top-level classes GeoNames defines. This chapter documents the extraction and sampling pipeline applied uniformly to twenty-one national extracts; Table 2.1 records the resulting counts for full transparency.

2.2 Sampling scheme

The central methodological decision in this volume is how to go from "all populated places in a country" (ranging from 909 for Libya to 557,994 for India) to a fixed-size sample suitable for pairwise-distance computation, without silently biasing the result toward or against any projection.

Two naive schemes were considered and rejected. Sampling only the largest cities by population would produce a point cloud dominated by capitals and a handful of major conurbations, understating the ordinary spatial spread of a country's settlement pattern and likely favouring whatever projection happens to handle the resulting sparse, corner-heavy configuration well. Sampling uniformly at random from all populated-place records, on the other hand, would under-represent genuinely major cities whenever they are numerically rare relative to the very large number of minor hamlets and suburbs also recorded by GeoNames (India's 557,994 populated-place records are overwhelmingly small villages).

Methodological Note

The scheme adopted blends both considerations, and is applied identically, with an identical fixed random seed, to every one of the twenty-one countries:

1. Sort all populated-place records by reported population, descending.
2. Take the top 40 (the country's genuine urban/administrative skeleton — for countries with fewer than 40 populated-place records, take all of them; no country in this study falls below 80 total, *cf.* Table 2.1).

Table 2.1: Data provenance: feature-class-P (populated place) record counts per country in the GeoNames extracts used, and the resulting $n = 80$ sample composition (top 40 by reported population, plus 40 more drawn by seeded random sampling from the remainder; countries with fewer than 80 total records would take all of them, but none does).

ISO	Country	Populated places (P)	Top-40	Seeded-40
IT	Italy	63,028	40	40
ES	Spain	30,894	40	40
PT	Portugal	16,614	40	40
HR	Croatia	11,631	40	40
SI	Slovenia	6,641	40	40
FR	France	81,598	40	40
DZ	Algeria	8,333	40	40
MA	Morocco	46,945	40	40
TN	Tunisia	2,245	40	40
LY	Libya	909	40	40
MR	Mauritania	7,653	40	40
AM	Armenia	2,275	40	40
AZ	Azerbaijan	5,295	40	40
GE	Georgia	5,419	40	40
TR	Turkey	54,607	40	40
IN	India	557,994	40	40
PK	Pakistan	150,394	40	40
AT	Austria	20,913	40	40
CH	Switzerland	13,199	40	40
GB	United Kingdom	43,712	40	40
IE	Ireland	12,159	40	40

3. From the remaining records, draw a uniform random sample (without replacement) of 40 more, using a deterministic seed derived from the ISO country code, so the exact sample is reproducible from Chapter 11 without needing to store point lists in this text.
4. The union, $n = 80$ points per country, is the sample used for every projection and every table in Chapters 3 to 7.

This is a stated, arbitrary-but-fixed choice, not tuned per country or per projection; it is applied before any projection is computed and never revisited afterward.

2.3 Reuse, unchanged, of the parent volume's IDI pipeline

Nothing about the definition of the Informational Distortion Index, or about the implementation of the twelve projections tested, is altered from Chapter 8 of the parent volume. For a sample $\mathcal{G} = \{p_1, \ldots, p_{80}\}$ of real populated places:

1. All $\binom{80}{2} = 3160$ pairwise geodesic distances are computed via the haversine formula on a spherical Earth of IUGG mean radius $R = 6371.0088$ km, exactly as in the parent volume's haversine formula.
2. Each of the twelve projections — Equirectangular, Mercator, Sinusoidal, Lambert Az-

imuthal Equal-Area (LAEA), Lambert Conformal Conic (LCC), Robinson, Aitoff, Winkel Tripel, Mollweide, Albers Equal-Area Conic, Hammer, and Van der Grinten — is applied to the same 80 points, and all 3160 pairwise planar Euclidean distances are computed from the projected coordinates.

3. The projected-distance distribution is rescaled by a single global constant (the median ratio of geodesic to projected pairwise distance) so that overall map scale, which carries no distortion information, does not contaminate the comparison, exactly as in the parent volume's Definition 8.1.

4. Both distributions are binned into 25 common histogram bins spanning $[0, \max(d_{\text{geo}})]$, and the Jensen–Shannon divergence between the two histograms, in bits, is the country's IDI value for that projection.

2.4 Projection centring for national-scale analysis

One genuine methodological adaptation is required, and is stated explicitly here because it is a substantive choice, not a cosmetic one. In the parent volume's global and regional-box grids, longitude was measured from the Greenwich meridian ($\lambda_0 = 0°$) throughout, which is a reasonable convention for whole-Earth or fixed-box analysis. For a per-country study, every projection is instead recentred on that country's own sample centroid (φ_0, λ_0) — the mean latitude and longitude of its 80 sampled points — before being evaluated. Concretely: for projections that expose an explicit centring parameter (LAEA, LCC, Aitoff, Mollweide, Albers, Hammer, Van der Grinten), the centroid is passed directly; for projections defined without one in the parent volume's code (Equirectangular, Mercator, Sinusoidal, Robinson, Winkel Tripel), longitudes are shifted by $\lambda - \lambda_0$ before the published formula is applied. Lambert Conformal Conic and Albers additionally use the sample's own 25^{th} and 75^{th} latitude percentiles as standard parallels, exactly as in the parent volume's Chapter 8.

Methodological Note

This recentring is necessary for the comparison to be meaningful at all: without it, a country far from the Greenwich meridian (India, centred near 78°E) would appear artificially distorted under every projection simply from being off-centre on a world map, confounding placement with the projection's intrinsic distortion structure — a different question from the one this volume asks. Recentring answers the question "if this projection were used as the natural, best-effort choice for mapping this specific country," which is the question relevant to an atlas maker choosing a projection for a national or regional map, and is also the fairer comparison across projections that differ in whether they expose a free centring parameter at all.

2.5 Aggregation across countries within a region

For each of the five regional groups defined in Chapter 1, a regional mean IDI per projection is computed as the arithmetic mean, over the countries in that region, of the country-level IDI values — not a re-pooling of all countries' raw point clouds into one larger sample. This choice keeps each country's contribution to the regional figure equal regardless of its population-record count (India's 557,994 populated-place records do not get more influence over the South Asia regional mean than Pakistan's 150,394 beyond what the $n = 80$ sampling already allows),

and it lets Chapters 3 to 7 report, transparently, both the per-country detail and the regional summary side by side.

2.6 Summary

Every subsequent numerical table in this volume follows from: (i) the fixed top-40-plus-seeded-40 sampling scheme of Section 2.2, (ii) the unmodified IDI definition and twelve projection formulas of the parent volume, (iii) the centroid-recentring convention of Section 2.4, and (iv) simple, unweighted, per-country averaging within each region. No step is projection-specific, country-specific, or adjusted after seeing a result.

Chapter 3

Mediterranean Europe: Italy, Spain, Portugal, Croatia, Slovenia, France

3.1 Regional character

The six Mediterranean European countries in this study span a wide range of territorial shapes at a broadly similar mid-latitude band ($36°$–$51°$N): the elongated, curved peninsula of Italy; the compact quadrilateral of Spain; the narrow north–south strip of Portugal; the crescent-shaped Adriatic coastline of Croatia; small, compact Slovenia; and France, the largest and most areally compact of the six. This diversity of shape, within one relatively narrow latitude band, makes the region a useful first test of whether the IDI ranking is driven mainly by latitude (a known effect — Mercator, in particular, is latitude-sensitive by construction) or is robust to shape variation at fixed latitude. It is also, by construction, the largest of the five regional groups in this volume, and the one whose constituent countries are individually best documented in GeoNames: Croatia's $11{,}631$ feature-class-P records and Slovenia's $6{,}641$ are the sparsest of the six, but both still comfortably exceed the $n = 80$ sample size by two orders of magnitude, leaving no country in this region anywhere near the kind of gazetteer thinness that becomes a genuine methodological concern in Chapter 4. Choosing this region to open the empirical survey is therefore a deliberate ordering decision: it lets the central finding be established first under close to ideal data conditions, before the harder cases — sparser gazetteers, more extreme elongation, higher latitude — are introduced one at a time in the chapters that follow.

3.2 Regional summary

3.3 Country-by-country detail

The regional mean and the headline first-place result do not exhaust what the full twelve-projection ranking has to say about each country individually. This section walks through each of the six rankings in turn, including the middle of the table — the six whole-Earth projections that are never seriously in contention for first place but whose relative order among themselves varies considerably from one country to the next — and reports, for each country, the size of the gap separating the last member of the regionally-fitted trio from the first whole-Earth projection to beat it. This gap is itself a new quantity, not previously tabulated in this chapter, computed directly from the same per-country rankings already given in Tables 3.2 to 3.7.

Italy (Table 3.2) gives the cleanest full ranking in the region: LAEA, Albers, and LCC occupy the top three exactly as expected, and the trio-to-field gap — the ratio between the weakest trio member (LCC, 4.62×10^{-5}) and the best whole-Earth projection (Mollweide, 1.78×10^{-4}) — is $3.85\times$, toward the narrower end of what this region produces. Mollweide's fourth-place finish, ahead of Sinusoidal (1.95×10^{-4}, fifth), is itself worth noting: in three of the other five

Table 3.1: Mediterranean Europe: regional mean Informational Distortion Index by projection, unweighted arithmetic mean over the 6 constituent countries (Italy, Spain, Portugal, Croatia, Slovenia, France).

Rank	Projection	Mean IDI (bits)
1	Albers Equal-Area Conic	2.76×10^{-5}
2	Lambert Azimuthal Equal-Area	3.11×10^{-5}
3	Lambert Conformal Conic	3.32×10^{-4}
4	Winkel Tripel	1.80×10^{-3}
5	Sinusoidal	1.81×10^{-3}
6	Mollweide	1.99×10^{-3}
7	Equirectangular	2.01×10^{-3}
8	Aitoff	3.41×10^{-3}
9	Robinson	3.90×10^{-3}
10	Van der Grinten	4.77×10^{-3}
11	Hammer	5.40×10^{-3}
12	Mercator	5.44×10^{-3}

countries in this region (Portugal, Croatia, Slovenia) Sinusoidal is the best-placed whole-Earth projection, so Italy is a partial exception to what would otherwise be a regional pattern. The middle of Italy's table — Winkel Tripel, Equirectangular — separates the fourth/fifth-place pair from a tightly clustered group of four projections (Aitoff, Hammer, Robinson, Van der Grinten, all between 1.87×10^{-3} and 2.15×10^{-3}) before Mercator closes the table at 2.80×10^{-3}, the single worst value recorded for any projection in Italy and, as already noted, worse than Mercator's showing in every other Mediterranean country examined here.

Spain (Table 3.3) is the region's one full-trio-sweep exception, and its complete ranking clarifies exactly how that exception arises: Winkel Tripel intrudes at third place (1.46×10^{-3}), ahead of LCC (1.83×10^{-3}, fourth), giving a comparatively modest trio-completion gap of $2.34\times$ once LCC is reached. What is easy to miss from the headline exception alone is that Sinusoidal — the projection responsible for the analogous exceptions in Algeria and India, discussed in Chapters 4 and 6 — is in fact one of the *weakest* whole-Earth performers for Spain, placing tenth of twelve at 9.08×10^{-3}. Spain's exception and Algeria's or India's exception are therefore not instances of the same underlying mechanism wearing different country labels; they are three genuinely distinct events, each driven by a different whole-Earth projection finding a different kind of local advantage. Spain's own worst-to-second-worst pair, Van der Grinten (2.08×10^{-2}) and Mercator (1.93×10^{-2}), are both markedly larger than any other value recorded anywhere in this region, making Spain, alongside Portugal, one of the two countries responsible for this region's most extreme whole-Earth failures.

Portugal (Table 3.4) is the region's other departure from the typical pattern, but in a different way: it is the only Mediterranean country where the two conic-family entries, Albers (9.10×10^{-6}) and LCC (4.69×10^{-5}), occupy first and second place ahead of LAEA (1.40×10^{-4}, third), rather than LAEA leading as it does in the other five countries of this region. This is consistent with Portugal's markedly elongated, narrow north–south shape, which plays directly to the strength of a projection with two independently placed standard parallels. The trio-completion gap here, $6.43\times$, is a full order of magnitude in the appropriate sense — Sinusoidal, fourth at 9.00×10^{-4}, is meaningfully closer to the trio than the rest of the table, which descends through Mollweide, Van der Grinten, Equirectangular, Mercator, Winkel Tripel, and Aitoff before Robinson and Hammer close out the ranking at 1.01×10^{-2} and 1.85×10^{-2} respectively, the latter the worst value recorded for Hammer in this region.

Table 3.2: Italy: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 63,028 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	1.92×10^{-6}
2	Albers Equal-Area Conic	1.99×10^{-5}
3	Lambert Conformal Conic	4.62×10^{-5}
4	Mollweide	1.78×10^{-4}
5	Sinusoidal	1.95×10^{-4}
6	Winkel Tripel	2.24×10^{-4}
7	Equirectangular	3.07×10^{-4}
8	Aitoff	1.87×10^{-3}
9	Hammer	1.98×10^{-3}
10	Robinson	2.00×10^{-3}
11	Van der Grinten	2.15×10^{-3}
12	Mercator	2.80×10^{-3}

Croatia (Table 3.5) returns the region's first exact-zero IDI (LAEA), with Albers (6.61×10^{-6}) and LCC (1.05×10^{-5}) close behind at the same resolution-floor scale already discussed for the Caucasus region in Chapter 5. The trio-completion gap, coincidentally also $6.43\times$ to fourth-place Sinusoidal (6.76×10^{-5}), matches Portugal's almost exactly despite the two countries having very different shapes, which is a reminder that this gap is not a simple monotonic function of elongation or compactness alone. Croatia's middle table — Winkel Tripel, Equirectangular, Mercator, Van der Grinten, Mollweide — is comparatively tightly packed between 1.23×10^{-4} and 5.55×10^{-4}, before Aitoff, Robinson, and Hammer occupy the bottom three ranks in that order, the same relative order Hammer and Robinson take in several other countries in this region.

Slovenia (Table 3.6) produces the narrowest trio-completion gap recorded anywhere in Mediterranean Europe, $1.75\times$, and the closest internal spacing within the trio itself: Albers (1.27×10^{-5}) and LCC (1.31×10^{-5}) differ by only a factor of 1.03, the tightest such pairing in the region. Sinusoidal, fourth at 2.30×10^{-5}, is close enough to the trio that a reader skimming only the top four ranks might mistake it for a fourth regionally-fitted contender; it is not, and the gap re-opens sharply further down the table, where Equirectangular and Mercator (fifth and sixth, both around 8.5×10^{-5}) separate it from Winkel Tripel, Van der Grinten, Mollweide, Aitoff, Robinson, and finally Hammer at 1.37×10^{-3}, the worst value recorded for Slovenia and, notably, smaller than the worst value recorded for every other country in this region — Slovenia's overall distortion landscape, at both ends of the ranking, is the mildest of the six.

France (Table 3.7) closes the region with its widest trio-completion gap, $10.47\times$, separating LCC (4.37×10^{-5}, third, itself very close to Albers's 4.11×10^{-5}) from Winkel Tripel (4.58×10^{-4}, fourth). France is also the one country in this region where Winkel Tripel, rather than Sinusoidal, is the best-placed whole-Earth projection, with Sinusoidal relegated to sixth place (5.89×10^{-4}) behind Equirectangular. The remainder of the table — Mollweide, Van der Grinten, Aitoff, Mercator, Robinson, Hammer — descends fairly evenly to a worst value of 2.75×10^{-3}, comparable in magnitude to Italy's worst value despite France's considerably larger and more compact territory, which is one more small illustration of the volume's general point that worst-case whole-Earth performance tracks country-specific shape rather than country size or gazetteer density.

Table 3.3: Spain: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 30,894 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	4.38×10^{-5}
2	Albers Equal-Area Conic	7.61×10^{-5}
3	Winkel Tripel	1.46×10^{-3}
4	Lambert Conformal Conic	1.83×10^{-3}
5	Aitoff	4.27×10^{-3}
6	Robinson	5.66×10^{-3}
7	Hammer	5.81×10^{-3}
8	Equirectangular	6.42×10^{-3}
9	Mollweide	7.55×10^{-3}
10	Sinusoidal	9.08×10^{-3}
11	Mercator	1.93×10^{-2}
12	Van der Grinten	2.08×10^{-2}

3.4 Interpretation

The regional mean (Table 3.1) places two of the three regionally-fitted projections, Albers (2.76×10^{-5}) and LAEA (3.11×10^{-5}), in the first two ranks, with LCC a clear third (3.32×10^{-4}) — itself still an order of magnitude below the best whole-Earth projection in this region, Winkel Tripel, at 1.80×10^{-3}, a gap of roughly $65\times$ between Albers and that best whole-Earth competitor. That the two equal-area regionally-fitted projections, Albers and LAEA, sit almost on top of each other at the regional-mean level, separated by little more than a factor of 1.1, while LCC trails by a further order of magnitude, is itself worth remarking on: it suggests that at this region's characteristic mix of shapes, the choice between an azimuthal and a conic equal-area construction matters far less than the choice of standard parallels does for the conformal conic, whose distortion behaviour is more sensitive to how well the two fitted parallels happen to bracket a given country's latitudinal spread. *First* place goes to a regionally-fitted projection in all six countries without exception (Tables 3.2 to 3.7): LAEA wins outright in Italy, Spain, Croatia, Slovenia, and France; Albers wins in Portugal, the sole exception to LAEA's otherwise unbroken run of first-place finishes across this region. The stronger claim — that the full trio always occupies the top three ranks — holds in five of the six countries but fails, honestly and specifically, in Spain: there, Winkel Tripel's IDI of 1.46×10^{-3} is lower than LCC's 1.83×10^{-3}, pushing LCC to fourth place. No mechanism internal to this study explains why a whole-Earth compromise projection should occasionally out-perform a regionally-fitted conic at a country's own latitude quartiles; it is reported here as an observed exception rather than reasoned away, and it recurs, in different guises, in Algeria (Chapter 4) and India (Chapter 6) as one of only three such exceptions across the entire twenty-one-country dataset.

Two further country-level details are worth flagging rather than folding silently into the regional average. First, Italy's curved peninsular shape — the case one might expect to most challenge a locally-fitted azimuthal or conic projection, since the territory bends appreciably in both latitude and longitude — still gives LAEA an IDI of 1.92×10^{-6}, an order of magnitude below Albers and more than three orders of magnitude below the worst whole-Earth projection tested for Italy (Mercator, 2.80×10^{-3}). The equal-area azimuthal centring absorbs the curvature well here: a single point-centred projection built around Italy's own sample centroid apparently loses very little accuracy to the fact that the peninsula itself curves away from

Table 3.4: Portugal: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 16,614 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Albers Equal-Area Conic	9.10×10^{-6}
2	Lambert Conformal Conic	4.69×10^{-5}
3	Lambert Azimuthal Equal-Area	1.40×10^{-4}
4	Sinusoidal	9.00×10^{-4}
5	Mollweide	2.09×10^{-3}
6	Van der Grinten	3.32×10^{-3}
7	Equirectangular	4.62×10^{-3}
8	Mercator	7.91×10^{-3}
9	Winkel Tripel	8.36×10^{-3}
10	Aitoff	9.81×10^{-3}
11	Robinson	1.01×10^{-2}
12	Hammer	1.85×10^{-2}

that centroid toward both its northern Alpine arc and its southern Sicilian extremity. Second, Mercator's rank varies sharply within this region: it is worst of all twelve projections for Italy, mid-table for Croatia and Slovenia (seventh and sixth respectively, comfortably ahead of Hammer, Robinson, and Aitoff in both cases), and among the three worst for Spain and France, consistent with Mercator's well-known latitude-dependent scale blow-up interacting with each country's specific north–south extent. Spain's worst projection is in fact Van der Grinten (2.08×10^{-2}), not Mercator (1.93×10^{-2}, second-worst), the only country in this region where Mercator is not the single worst performer — a small but genuine exception to what would otherwise be a clean regional pattern, and one more instance of the volume's general finding that the identity of the *worst* projection is considerably less stable across countries than the identity of the best one.

Portugal deserves a further word precisely because it is the region's one departure from LAEA supremacy at first place. Portugal's territory is a narrow north–south coastal strip, elongated along almost exactly the axis that a conic projection's twin standard parallels are best equipped to correct for, and Albers's win there (9.10×10^{-6}, with LCC a close second at 4.69×10^{-5} and LAEA a more distant third at 1.40×10^{-4}) is consistent with the same shape-dependent mechanism explored quantitatively, and found only weakly supported, in Section 9.4 of Chapter 9: an elongated territory tends to favour a projection with a directional correction built in over one that corrects distortion symmetrically around a single centre. Portugal's case is one of the cleanest illustrations of that intuition anywhere in the volume, even though the formal correlation test across all twenty-one countries does not reach statistical significance.

3.5 Summary

Across six countries of markedly different shape at similar latitude, the headline finding is unanimous at first place: a projection fitted to the region, using only information available at the time of mapmaking (the region's own centroid and latitude quartiles), wins outright in every country, typically by one to three orders of magnitude over the best whole-Earth alternative, regardless of whether the territory is compact (Slovenia), elongated (Portugal), or curved (Italy). The one qualification, reported rather than smoothed over, is Spain, where Winkel

Table 3.5: Croatia: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 11,631 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	0
2	Albers Equal-Area Conic	6.61×10^{-6}
3	Lambert Conformal Conic	1.05×10^{-5}
4	Sinusoidal	6.76×10^{-5}
5	Winkel Tripel	1.23×10^{-4}
6	Equirectangular	1.53×10^{-4}
7	Mercator	4.39×10^{-4}
8	Van der Grinten	5.19×10^{-4}
9	Mollweide	5.55×10^{-4}
10	Aitoff	1.58×10^{-3}
11	Robinson	1.82×10^{-3}
12	Hammer	2.01×10^{-3}

Tripel edges out LCC for third place — a reminder that the regionally-fitted trio's dominance is strongest and most consistent at the top of the ranking, not uniformly through every rank below it. Taken together with the near-parity between Albers and LAEA at the regional-mean level, and with Portugal's single clean departure from LAEA's otherwise unbroken first-place run, Mediterranean Europe establishes the volume's central pattern under close to the most favourable data conditions available in this study, setting a baseline against which the sparser, more extreme, or higher-latitude regions of Chapters 4 to 7 can be usefully compared.

Table 3.6: Slovenia: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 6,641 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	0
2	Albers Equal-Area Conic	1.27×10^{-5}
3	Lambert Conformal Conic	1.31×10^{-5}
4	Sinusoidal	2.30×10^{-5}
5	Equirectangular	8.36×10^{-5}
6	Mercator	8.67×10^{-5}
7	Winkel Tripel	1.73×10^{-4}
8	Van der Grinten	4.67×10^{-4}
9	Mollweide	4.85×10^{-4}
10	Aitoff	1.05×10^{-3}
11	Robinson	1.18×10^{-3}
12	Hammer	1.37×10^{-3}

Table 3.7: France: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 81,598 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	8.53×10^{-7}
2	Albers Equal-Area Conic	4.11×10^{-5}
3	Lambert Conformal Conic	4.37×10^{-5}
4	Winkel Tripel	4.58×10^{-4}
5	Equirectangular	5.05×10^{-4}
6	Sinusoidal	5.89×10^{-4}
7	Mollweide	1.05×10^{-3}
8	Van der Grinten	1.34×10^{-3}
9	Aitoff	1.90×10^{-3}
10	Mercator	2.11×10^{-3}
11	Robinson	2.66×10^{-3}
12	Hammer	2.75×10^{-3}

Chapter 4

The Maghreb: Algeria, Morocco, Tunisia, Libya, Mauritania

4.1 Regional character

The five Maghrebi and western Saharan countries in this study present the sharpest test of settlement-driven sampling in the entire volume: unlike the compact, densely gazetteered countries of Mediterranean Europe, several of these territories are vast, sparsely populated, and unevenly documented in GeoNames. Libya contributes only 909 populated-place records nationwide (Table 2.1), and Tunisia $2{,}245$ — both still comfortably above the $n = 80$ sample size, but leaving markedly less redundancy in the "seeded-random" half of the sample than, say, France's $81{,}598$ records. Algeria, Morocco, and Mauritania are each geographically enormous relative to their coastal population concentration, meaning the sampled settlements are overwhelmingly clustered along a relatively narrow coastal or riverine strip rather than spread uniformly across the nominal national bounding box — itself a genuine empirical fact about these countries' human geography, not a sampling artefact to be corrected away. This makes the Maghreb, in effect, a stress test of the sampling scheme fixed in Section 2.2: if the top-40-plus-seeded-40 construction were quietly biasing the result in some region-specific way, this is the region where the thinness of the underlying gazetteer and the extremity of the coastal-versus-interior population imbalance would be most likely to expose it.

4.2 Regional summary

4.3 Country-by-country detail

As in Chapter 3, the regional mean and the first-place headline understate how much the five full twelve-projection rankings actually differ from one another in their middle and lower ranks. This section works through each country's complete ranking and reports the trio-completion gap for each, computed the same way as in Section 3.3.

Algeria (Table 4.2) is the region's sole trio-sweep exception, and its full ranking shows the exception is not marginal: Sinusoidal takes second place outright (1.63×10^{-4}), ahead of both LCC (1.98×10^{-4}, third) and Albers (2.74×10^{-4}, fourth), so it is Albers, not LCC, that gets displaced from the trio's usual top-three slot here. The trio-completion gap, measured from Albers to fifth-place Equirectangular (6.58×10^{-4}), is $2.40\times$, comparable to Spain's exception gap in Chapter 3. Algeria's lower table is unusually spread out even by this volume's standards: Robinson (sixth, 1.95×10^{-3}) sits nearly three times below Winkel Tripel (seventh, 5.16×10^{-3}), and the bottom four — Mollweide, Mercator, Aitoff, Van der Grinten, Hammer — are all clustered within a comparatively narrow band around 1.1–1.2×10^{-2}, with Hammer worst at 1.22×10^{-2}, the largest value recorded anywhere in this region.

Table 4.1: Maghreb: regional mean Informational Distortion Index by projection, unweighted arithmetic mean over the 5 constituent countries (Algeria, Morocco, Tunisia, Libya, Mauritania).

Rank	Projection	Mean IDI (bits)
1	Lambert Azimuthal Equal-Area	1.99×10^{-5}
2	Albers Equal-Area Conic	6.77×10^{-5}
3	Lambert Conformal Conic	1.19×10^{-4}
4	Sinusoidal	2.78×10^{-4}
5	Equirectangular	5.44×10^{-4}
6	Robinson	1.02×10^{-3}
7	Mollweide	3.22×10^{-3}
8	Winkel Tripel	3.26×10^{-3}
9	Mercator	3.31×10^{-3}
10	Aitoff	3.54×10^{-3}
11	Van der Grinten	3.83×10^{-3}
12	Hammer	4.74×10^{-3}

Morocco (Table 4.3) gives the smallest non-zero LAEA value recorded anywhere in the entire twenty-one-country study, 4.59×10^{-7}, ahead of Albers (1.17×10^{-5}) and LCC (1.41×10^{-5}) by more than an order of magnitude, and the trio-completion gap to fourth-place Sinusoidal (8.31×10^{-5}) is $5.91\times$. Morocco's full ranking is also the source of this region's most striking worst-place result: Winkel Tripel, which performs respectably in several other regions (it is the best-placed whole-Earth projection for France in Chapter 3), is the single worst projection for Morocco at 7.19×10^{-4}, trailing even Hammer (eleventh, 4.14×10^{-4}) and Mollweide (tenth, 3.79×10^{-4}). No other country in the study places Winkel Tripel last, which underscores the volume's general point, made repeatedly across Chapters 3, 5 and 7, that a whole-Earth projection's relative failure is a function of the specific country's shape and latitude rather than a fixed property of the projection itself.

Tunisia (Table 4.4) returns an exact-zero LAEA value, joining Croatia, Slovenia, Armenia, Georgia, and several Atlantic-region countries in the resolution-floor cluster discussed at length in Chapters 5 and 7. Albers (3.04×10^{-6}) and LCC (2.35×10^{-5}) complete the trio, and the trio-completion gap to Sinusoidal (2.05×10^{-4}, fourth) is $8.72\times$, among the wider gaps recorded in this region. Tunisia's table is otherwise unremarkable in its middle ranks — Equirectangular, Robinson, Mercator, Van der Grinten descend in a fairly regular sequence — before Aitoff and Mollweide arrive almost tied (3.29×10^{-3} and 3.31×10^{-3}) ahead of Winkel Tripel and, worst of all, Hammer at 7.09×10^{-3}.

Libya (Table 4.5), with the sparsest gazetteer of any country in the study at 909 feature-class-P records, nonetheless produces a completely ordinary trio-sweep result: LAEA (2.22×10^{-5}) and Albers (2.44×10^{-5}) are close enough to be read almost as a tie (ratio 1.10), with LCC a clear third at 3.32×10^{-4} and a trio-completion gap to Sinusoidal (fourth, 8.94×10^{-4}) of $2.69\times$. Libya's bottom of the table is where its data sparsity, if it were going to show through as instability, would be most likely to appear, and to some extent it does: the last five projections (Aitoff, Mercator, Hammer, Van der Grinten, Mollweide) are all packed within a comparatively narrow band between 3.05×10^{-3} and 4.79×10^{-3}, a tighter bottom-of-table clustering than in Algeria or Tunisia, consistent with a smaller, less redundant sample giving the weakest projections less room to separate from one another. Mollweide's last-place finish here, at 4.79×10^{-3}, is nonetheless barely distinguishable from Winkel Tripel's second-worst 4.76×10^{-3} immediately above it.

Table 4.2: Algeria: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 8,333 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	6.73×10^{-5}
2	Sinusoidal	1.63×10^{-4}
3	Lambert Conformal Conic	1.98×10^{-4}
4	Albers Equal-Area Conic	2.74×10^{-4}
5	Equirectangular	6.58×10^{-4}
6	Robinson	1.95×10^{-3}
7	Winkel Tripel	5.16×10^{-3}
8	Mollweide	5.73×10^{-3}
9	Mercator	1.08×10^{-2}
10	Aitoff	1.10×10^{-2}
11	Van der Grinten	1.18×10^{-2}
12	Hammer	1.22×10^{-2}

Mauritania (Table 4.6), already discussed in Section 4.3 above for its distinctive coastal-versus-interior sample composition, produces the narrowest trio-completion gap in this region at $1.65\times$ (LCC, 2.68×10^{-5}, to Sinusoidal, 4.41×10^{-5}). Its full ranking is otherwise notable for placing Aitoff and Hammer (fifth and sixth, 9.38×10^{-5} and 9.96×10^{-5}) well ahead of Equirectangular, Mercator, and Van der Grinten (seventh through ninth), an ordering not seen for these particular projections anywhere else in the Maghreb, before Robinson, Winkel Tripel, and Mollweide close out the table — Mollweide worst, as in Libya, at 1.88×10^{-3}.

4.4 Interpretation

The same qualitative pattern as Mediterranean Europe reappears here despite the very different settlement geometry: LAEA leads the regional mean (Table 4.1, 1.99×10^{-5}), with Albers second (6.77×10^{-5}) and LCC third (1.19×10^{-4}), all three at least an order of magnitude below the best whole-Earth projection, Sinusoidal (2.78×10^{-4}). The gap between LAEA and Albers here, a factor of roughly 3.4, is noticeably wider than the near-parity the same two projections showed in Mediterranean Europe (a factor of only 1.1), which is consistent with a region whose constituent countries are individually more elongated or more irregularly shaped than the Mediterranean set: an azimuthal projection recentred on a single point absorbs that kind of irregularity somewhat better, on average, than a conic projection whose two standard parallels are fixed by latitude quartiles alone and take no account of longitudinal spread. LAEA in fact takes first place in all five countries individually (Tables 4.2 to 4.6) — unlike Mediterranean Europe, no country in this region reassigns first place to LCC or Albers, making the Maghreb the only region in the study, alongside South Asia, where a single projection sweeps first place in every one of its constituent countries.

Algeria is the one country in this region (and, together with Spain and India, one of three in the entire study) where the fitted trio does not sweep the top three outright: Sinusoidal's IDI of 1.63×10^{-4} beats Albers's 2.74×10^{-4}, pushing Albers to fourth place behind LAEA, Sinusoidal, and LCC. Algeria's settlement pattern — overwhelmingly coastal, with the sampled interior points drawn from a vast, sparsely populated Saharan expanse — plausibly gives a true-scale, meridian-preserving projection like Sinusoidal an unusual advantage here that a

Table 4.3: Morocco: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 46,945 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	4.59×10^{-7}
2	Albers Equal-Area Conic	1.17×10^{-5}
3	Lambert Conformal Conic	1.41×10^{-5}
4	Sinusoidal	8.31×10^{-5}
5	Robinson	1.67×10^{-4}
6	Equirectangular	1.96×10^{-4}
7	Van der Grinten	2.45×10^{-4}
8	Aitoff	2.58×10^{-4}
9	Mercator	2.82×10^{-4}
10	Mollweide	3.79×10^{-4}
11	Hammer	4.14×10^{-4}
12	Winkel Tripel	7.19×10^{-4}

conic projection's twin standard parallels do not confer: Sinusoidal preserves area and true scale along every parallel, which may partially compensate for a settlement pattern concentrated at one latitude band (the Mediterranean coast) with only a thin scatter of interior points pulling the sample's latitude quartiles toward the Sahara. This is offered as a plausible mechanism rather than a proven one; no analytic derivation in this volume isolates Sinusoidal's specific advantage for a coastal-plus-sparse-interior settlement geometry, and the honest position is that Algeria's exception, like Spain's and India's, is reported because it occurs, not because its cause has been established with certainty.

The worst-performing projection is more variable here than in Mediterranean Europe: Hammer is worst for Algeria and Tunisia, Winkel Tripel for Morocco, and Mollweide for both Libya and Mauritania. This variability among the *losing* projections, contrasted with the near-total stability of first place, is itself an informative pattern: whole-Earth compromise projections do not fail for a single, shared reason at national scale — each trades off global properties (equal-area pseudocylindrical shape for Mollweide and Hammer, a fixed standard parallel for Winkel Tripel) that interact differently with each country's specific latitude and elongation. Regionally-fitted projections, by contrast, succeed for the same reason in every case: they are re-optimized, by construction, using exactly the country's own centroid and latitude spread. It is worth noting, too, that Morocco returns the smallest LAEA value recorded anywhere in this region (4.59×10^{-7}, an order of magnitude below even Algeria's or Mauritania's LAEA figures), despite Morocco having by far the largest gazetteer of the five Maghrebi countries (46,945 records) — a reminder that gazetteer size and IDI magnitude are not obviously related once the fixed $n = 80$ sampling scheme has already been applied; a densely documented country does not automatically produce a smaller or larger IDI than a sparsely documented one, since both are reduced to the same 80-point sample before any projection is evaluated.

Methodological Note

Mauritania is worth a specific methodological remark. Its interior is dominated, in the GeoNames extract, by sparse desert settlements, so the "top 40 by population" half of the sample is disproportionately coastal and riverine (concentrated near Nouakchott and the Senegal River valley), while the "seeded 40" half draws more evenly from the interior.

Table 4.4: Tunisia: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 2,245 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	0
2	Albers Equal-Area Conic	3.04×10^{-6}
3	Lambert Conformal Conic	2.35×10^{-5}
4	Sinusoidal	2.05×10^{-4}
5	Equirectangular	2.48×10^{-4}
6	Robinson	6.14×10^{-4}
7	Mercator	1.46×10^{-3}
8	Van der Grinten	2.43×10^{-3}
9	Aitoff	3.29×10^{-3}
10	Mollweide	3.31×10^{-3}
11	Winkel Tripel	3.80×10^{-3}
12	Hammer	7.09×10^{-3}

Table 4.5: Libya: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 909 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	2.22×10^{-5}
2	Albers Equal-Area Conic	2.44×10^{-5}
3	Lambert Conformal Conic	3.32×10^{-4}
4	Sinusoidal	8.94×10^{-4}
5	Robinson	1.05×10^{-3}
6	Equirectangular	1.46×10^{-3}
7	Aitoff	3.05×10^{-3}
8	Mercator	3.79×10^{-3}
9	Hammer	3.89×10^{-3}
10	Van der Grinten	4.42×10^{-3}
11	Winkel Tripel	4.76×10^{-3}
12	Mollweide	4.79×10^{-3}

This blend, rather than either extreme alone, is precisely the situation the two-part sampling scheme of Section 2.2 was designed to avoid distorting: neither an all-capital-cities sample nor a naively uniform one would have represented Mauritania's actual settlement geography as honestly. Mauritania's own result — LAEA first at 9.79×10^{-6}, with the full trio sweeping the top three and Sinusoidal a comparatively distant fourth at 4.41×10^{-5} — does not show any obvious sign of having been distorted by this coastal/interior blend, but the blend's existence is disclosed here rather than left for the reader to infer.

4.5 Summary

Despite the widest variation in territorial size, coastal concentration, and gazetteer density of any region in this study, the Maghreb reproduces the central finding without exception:

Table 4.6: Mauritania: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 7,653 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	9.79×10^{-6}
2	Albers Equal-Area Conic	2.48×10^{-5}
3	Lambert Conformal Conic	2.68×10^{-5}
4	Sinusoidal	4.41×10^{-5}
5	Aitoff	9.38×10^{-5}
6	Hammer	9.96×10^{-5}
7	Equirectangular	1.54×10^{-4}
8	Mercator	1.92×10^{-4}
9	Van der Grinten	2.04×10^{-4}
10	Robinson	1.31×10^{-3}
11	Winkel Tripel	1.83×10^{-3}
12	Mollweide	1.88×10^{-3}

regionally-fitted projections dominate whole-Earth projections by one to two orders of magnitude in every one of the five countries. LAEA's clean sweep of first place across all five countries, unmatched among the trio elsewhere in the volume outside South Asia, and the single, mechanistically plausible exception in Algeria's trio-sweep, together suggest that the coastal-versus-interior population imbalance so distinctive of this region's human geography does not, in itself, threaten the volume's central claim — if anything, the wider gap observed here between LAEA and its nearest regionally-fitted competitor indicates that shape irregularity of the kind common across the Maghreb favours the azimuthal construction over the conic one somewhat more decisively than the milder shapes of Mediterranean Europe did.

Chapter 5

The Caucasus and Anatolia: Armenia, Azerbaijan, Georgia, Turkey

5.1 Regional character

These four countries share a mountainous, mid-latitude (37°–42°N) setting at the juncture of Europe and Asia, but differ sharply in scale: Armenia, Azerbaijan, and Georgia are each small, broadly compact territories, while Turkey is an order of magnitude larger in area and considerably more elongated east–west. This region therefore tests the same question as Chapter 3 (does shape variation at fixed latitude disturb the ranking?) but with a starker size contrast, and with two of its four countries (Armenia and Azerbaijan) among the sparser gazetteer entries in the study (Table 2.1: 2,275 and 5,295 feature-class-P records respectively, smaller than every Mediterranean European country and every Maghrebi country except Libya and Tunisia). The pairing of three small, compact countries against one large, elongated one within a single regional group also makes this the region best suited, within the volume's five-way regional split, to examine directly whether the trio's internal ordering — which of LAEA, Albers, or LCC takes first place — responds to elongation at all, a question taken up more systematically, and with a more equivocal answer, in Section 9.4 of Chapter 9.

5.2 Regional summary

Table 5.1: Caucasus and Anatolia: regional mean Informational Distortion Index by projection, unweighted arithmetic mean over the 4 constituent countries (Armenia, Azerbaijan, Georgia, Turkey).

Rank	Projection	Mean IDI (bits)
1	Lambert Azimuthal Equal-Area	2.97×10^{-6}
2	Lambert Conformal Conic	4.59×10^{-5}
3	Albers Equal-Area Conic	4.60×10^{-5}
4	Mollweide	1.75×10^{-4}
5	Sinusoidal	2.50×10^{-4}
6	Equirectangular	3.17×10^{-4}
7	Winkel Tripel	4.42×10^{-4}
8	Mercator	4.47×10^{-4}
9	Van der Grinten	7.73×10^{-4}
10	Robinson	9.20×10^{-4}
11	Aitoff	1.07×10^{-3}
12	Hammer	1.95×10^{-3}

Table 5.2: Armenia: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 2,275 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	0
2	Albers Equal-Area Conic	8.77×10^{-6}
3	Lambert Conformal Conic	1.00×10^{-5}
4	Mollweide	6.41×10^{-5}
5	Equirectangular	1.24×10^{-4}
6	Sinusoidal	1.32×10^{-4}
7	Mercator	2.93×10^{-4}
8	Winkel Tripel	3.64×10^{-4}
9	Van der Grinten	6.04×10^{-4}
10	Robinson	6.35×10^{-4}
11	Aitoff	7.58×10^{-4}
12	Hammer	1.48×10^{-3}

5.3 Country-by-country detail

As in Chapters 3 and 4, the regional mean and the headline first-place finding leave most of each country's full twelve-projection ranking undiscussed. This section works through each of the four rankings in turn and reports the trio-completion gap for each, computed the same way as in Sections 3.3 and 4.3.

Armenia (Table 5.2) gives the region's first exact-zero LAEA result, with Albers (8.77×10^{-6}) and LCC (1.00×10^{-5}) close behind, within a factor of 1.14 of each other. The trio-completion gap, from LCC to fourth-place Mollweide (6.41×10^{-5}), is $6.41\times$. Sinusoidal, which intrudes into the top four in several other regions, is only sixth here (1.32×10^{-4}), behind Equirectangular (fifth, 1.24×10^{-4}) — Armenia is one of the clearer cases in the volume where Sinusoidal shows no special local advantage at all. The lower table descends fairly evenly through Mercator, Winkel Tripel, Van der Grinten, and Robinson to Aitoff (eleventh, 7.58×10^{-4}) and, worst of all, Hammer at 1.48×10^{-3}.

Azerbaijan (Table 5.3) records the second-smallest non-zero LAEA value in the entire study, 9.30×10^{-7}, trailing only Morocco's; LCC (4.74×10^{-6}) edges Albers (5.14×10^{-6}) for second place, the only Caucasus country where LCC rather than Albers takes the trio's middle rank. The trio-completion gap, from Albers to fourth-place Sinusoidal (8.88×10^{-5}), is $17.28\times$ — the widest recorded anywhere in this region. Sinusoidal's fourth place here is still nowhere near displacing a trio member, unlike its behaviour in Algeria or India (Chapters 4 and 6). The remainder of the table is comparatively regular, descending through Mollweide, Equirectangular, Van der Grinten, Mercator, and Winkel Tripel before Robinson, Aitoff, and Hammer close the ranking at 1.03×10^{-3}, 1.49×10^{-3}, and 2.32×10^{-3} respectively — the largest worst-case value recorded for any Caucasus country.

Georgia (Table 5.4) returns the region's second exact-zero LAEA value and, as already noted, a tie between LCC and Albers at 1.5825×10^{-4} to the precision reported. A second, independent tie occurs lower in the same table: Mollweide (fifth) and Winkel Tripel (sixth) both register 2.21×10^{-4}, making Georgia the only country in the volume with two separate ties among its twelve projections. The trio-completion gap, from the tied Albers/LCC pair to fourth-place Equirectangular (2.17×10^{-4}), is $1.37\times$ — the narrowest in this region. Sinusoidal is seventh

Table 5.3: Azerbaijan: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 5,295 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	9.30×10^{-7}
2	Lambert Conformal Conic	4.74×10^{-6}
3	Albers Equal-Area Conic	5.14×10^{-6}
4	Sinusoidal	8.88×10^{-5}
5	Mollweide	1.18×10^{-4}
6	Equirectangular	1.71×10^{-4}
7	Van der Grinten	3.15×10^{-4}
8	Mercator	3.68×10^{-4}
9	Winkel Tripel	5.71×10^{-4}
10	Robinson	1.03×10^{-3}
11	Aitoff	1.49×10^{-3}
12	Hammer	2.32×10^{-3}

here (2.30×10^{-4}), immediately behind the Mollweide/Winkel Tripel tie, before Mercator, Van der Grinten, Aitoff, and Robinson descend to Hammer's last-place 7.81×10^{-4}, itself the smallest worst-case value recorded for any Caucasus country.

Turkey (Table 5.5) is the region's, and one of only two countries in the entire study's, LCC-over-LAEA exception already discussed above, with all three trio members clustered within a factor of 1.1 of one another. Despite that tight three-way clustering at the top, the trio-completion gap to fourth-place Mollweide (2.98×10^{-4}) is $25.25\times$ — the widest recorded anywhere in the Caucasus, and a reminder that a narrow spread within the trio and a wide gap to the rest of the field are independent properties of a ranking rather than opposite ends of the same scale. This is also the region's largest gazetteer by a wide margin ($54{,}607$ feature-class-P records, roughly ten times Georgia's or Azerbaijan's), yet it produces the region's widest rather than narrowest trio-completion gap, one more instance of the volume's recurring point that gazetteer size does not obviously predict any property of the resulting ranking. The middle table — Sinusoidal, Winkel Tripel, Equirectangular, Mercator — descends fairly evenly before Robinson, Aitoff, and Van der Grinten cluster closely together (between 1.53×10^{-3} and 1.76×10^{-3}) ahead of Hammer's last-place 3.24×10^{-3}, the largest worst-case value recorded anywhere in this region.

5.4 Interpretation

Several of the lowest IDI values recorded anywhere in the entire study occur within this region: Armenia's and Georgia's LAEA values are each exactly 0 (Tables 5.2 and 5.4), and Azerbaijan's LAEA value, while not exactly zero, is the second-smallest non-zero LAEA figure recorded anywhere in the study at 9.30×10^{-7} (Table 5.3), trailing only Morocco's 4.59×10^{-7} (Chapter 4) for the single smallest non-zero value in the entire twenty-one-country dataset. This is not a numerical coincidence peculiar to this region so much as a consequence of country size interacting with the fixed 25-bin histogram used throughout: for a small, compact territory, the projected-distance histogram for a well-fitted projection can coincide bin-for-bin with the geodesic-distance histogram at this resolution, driving the JSD computationally to zero even though the underlying continuous distributions are, in principle, not identical. Section 9.1 of

Table 5.4: Georgia: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 5,419 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	0
2	Lambert Conformal Conic	1.58×10^{-4}
3	Albers Equal-Area Conic	1.58×10^{-4}
4	Equirectangular	2.17×10^{-4}
5	Mollweide	2.21×10^{-4}
6	Winkel Tripel	2.21×10^{-4}
7	Sinusoidal	2.30×10^{-4}
8	Mercator	2.67×10^{-4}
9	Van der Grinten	4.17×10^{-4}
10	Aitoff	4.53×10^{-4}
11	Robinson	4.82×10^{-4}
12	Hammer	7.81×10^{-4}

Chapter 9 documents this bin-count sensitivity directly by rerunning the pipeline at coarser and finer resolutions; it reappears here as expected, and is flagged rather than treated as a more remarkable finding than it is. Georgia additionally shows LCC and Albers tied, to the precision reported, at 1.5825×10^{-4} each — a coincidence of the same fixed-resolution binning rather than a claim that the two projections are analytically identical for this country; at a different bin count, or with more decimal places reported, the tie would almost certainly resolve into an ordering, however narrow.

Turkey, the outlier in size and shape, is one of only two countries in the whole study (with the Mediterranean's Spain and the Maghreb's Algeria showing the trio itself imperfectly ordered, though not displaced from the top) where LCC edges out LAEA for first place, and by the narrowest margin recorded anywhere in the entire twenty-one-country study: 1.074×10^{-5} against 1.094×10^{-5}, a ratio of just 1.019. This is consistent with the pattern already seen: Turkey's territory spans a considerably wider range of longitude than latitude, an elongation unmatched by any other single country in the study, and it is exactly this kind of shape that should, on general grounds, favour a conic projection's directional standard-parallel construction over an azimuthal projection's rotationally symmetric one. Turkey's case is, in fact, the single clearest anecdotal support for the elongation–margin hypothesis examined quantitatively across all twenty-one countries in Section 9.4, where it anchors one extreme of a correlation that, taken as a whole, does not reach statistical significance; the honesty of that later chapter's negative result is only strengthened by acknowledging here, plainly, that Turkey alone looks exactly like the hypothesis predicts. Turkey's third-ranked Albers value, 1.18×10^{-5}, is close enough to both LCC and LAEA that all three regionally-fitted projections sit within a factor of 1.1 of one another — the tightest three-way clustering at the top of any country's ranking table in the volume.

The worst-performing projection is Hammer in all four countries without exception — the only region in the study where a single projection occupies last place across every constituent country. Aitoff is second-worst in Armenia and Azerbaijan and third-worst in Georgia and Turkey, suggesting that the Aitoff/Hammer family of azimuthal projections derived from halving the longitude is specifically ill-suited to this region's combination of latitude and shape, though establishing why with the same rigour as the Chebyshev/Milnor argument of Chapter 3 in the parent volume would require the analytic treatment this volume has explicitly declined

Table 5.5: Turkey: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 54,607 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Conformal Conic	1.07×10^{-5}
2	Lambert Azimuthal Equal-Area	1.09×10^{-5}
3	Albers Equal-Area Conic	1.18×10^{-5}
4	Mollweide	2.98×10^{-4}
5	Sinusoidal	5.50×10^{-4}
6	Winkel Tripel	6.12×10^{-4}
7	Equirectangular	7.57×10^{-4}
8	Mercator	8.59×10^{-4}
9	Robinson	1.53×10^{-3}
10	Aitoff	1.59×10^{-3}
11	Van der Grinten	1.76×10^{-3}
12	Hammer	3.24×10^{-3}

to force (Chapter 10). It is worth noting that the region's worst-case IDI values are also, on the whole, more moderate than those recorded elsewhere in the study: Hammer's value never exceeds 3.24×10^{-3} (Turkey) in this region, well below the worst values recorded in Mediterranean Europe (Spain's Van der Grinten at 2.08×10^{-2}) or the Maghreb (Algeria's Hammer at 1.22×10^{-2}), which suggests that the mountainous, mid-latitude character shared by all four countries here produces a somewhat gentler overall distortion landscape for whole-Earth projections than the more latitude- or shape-extreme regions elsewhere in the volume, even as the regionally-fitted trio's advantage over that landscape remains just as decisive in relative terms.

5.5 Summary

The Caucasus and Anatolia region reproduces the central finding at its most extreme: among the lowest IDI values recorded anywhere in the study (two exact zeros and a near-zero, a resolution artefact of the fixed bin count rather than a claim of perfect fidelity), the narrowest LAEA/LCC margin recorded anywhere (Turkey, a ratio of 1.019), and the only region-wide unanimous last-place projection (Hammer) in the entire volume. Turkey's status as the sole country in the study where an elongated, conic-favouring shape and a genuine LCC win coincide gives this region an outsized role in the elongation discussion of Section 9.4, even though the formal twenty-one-country test conducted there does not, on its own, establish the relationship Turkey's individual case illustrates so cleanly.

Chapter 6

South Asia: India and Pakistan

6.1 Regional character

India and Pakistan are, between them, the two most heavily gazetteered countries in this entire study by a wide margin: India alone contributes 557,994 populated-place records, and Pakistan 150,394 (Table 2.1), together outnumbering the combined populated-place count of all nineteen other countries in the study by a considerable margin. This region is also the only one consisting of just two countries, so no regional mean is computed over more than two values; the regional summary table below should be read with that caveat, though the qualitative pattern it reports is, as elsewhere, confirmed independently in both underlying per-country tables. The pairing is a deliberate contrast in scale as much as a geographic grouping: India's territory spans one of the widest latitudinal ranges of any single country in this study, while Pakistan, though itself substantial, is markedly more compact in its north–south extent, giving this two-country region an internal size and shape contrast not unlike the Armenia/Georgia-versus-Turkey pairing examined in Chapter 5, but concentrated here into just two data points rather than four.

6.2 Regional summary

Table 6.1: South Asia: regional mean Informational Distortion Index by projection, unweighted arithmetic mean over the 2 constituent countries (India, Pakistan).

Rank	Projection	Mean IDI (bits)
1	Lambert Azimuthal Equal-Area	1.00×10^{-5}
2	Albers Equal-Area Conic	6.13×10^{-5}
3	Lambert Conformal Conic	8.04×10^{-5}
4	Sinusoidal	9.82×10^{-5}
5	Equirectangular	1.71×10^{-4}
6	Aitoff	4.15×10^{-4}
7	Hammer	4.63×10^{-4}
8	Robinson	6.56×10^{-4}
9	Van der Grinten	6.86×10^{-4}
10	Mercator	8.39×10^{-4}
11	Winkel Tripel	1.64×10^{-3}
12	Mollweide	1.76×10^{-3}

Table 6.2: India: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 557,994 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	1.73×10^{-5}
2	Sinusoidal	1.10×10^{-4}
3	Albers Equal-Area Conic	1.14×10^{-4}
4	Lambert Conformal Conic	1.43×10^{-4}
5	Equirectangular	2.06×10^{-4}
6	Aitoff	3.78×10^{-4}
7	Hammer	4.01×10^{-4}
8	Van der Grinten	8.13×10^{-4}
9	Mercator	1.11×10^{-3}
10	Robinson	1.21×10^{-3}
11	Winkel Tripel	2.32×10^{-3}
12	Mollweide	2.45×10^{-3}

6.3 Country-by-country detail

As in the preceding regional chapters, this section reports the full ranking and the trio-completion gap for each of the two countries in this region.

India (Table 6.2) is the volume's most heavily documented trio-sweep exception: Sinusoidal's second-place finish (1.10×10^{-4}) displaces LCC to fourth (1.43×10^{-4}), with Albers holding third (1.14×10^{-4}) by a narrow margin over Sinusoidal itself, a factor of only 1.04 separating second from third place. The trio-completion gap, measured from LCC to fifth-place Equirectangular (2.06×10^{-4}), is $1.44\times$ — one of the narrowest recorded anywhere in the volume, consistent with Sinusoidal having already absorbed most of the separation that would otherwise appear between the trio and the rest of the field. The remaining table is comparatively compressed toward the bottom: Aitoff and Hammer (sixth and seventh, 3.78×10^{-4} and 4.01×10^{-4}) sit closer together than the corresponding pair in most other countries, before Van der Grinten, Mercator, and Robinson descend to Winkel Tripel (eleventh, 2.32×10^{-3}) and, worst of all, Mollweide at 2.45×10^{-3}.

Pakistan (Table 6.3) gives the region's clean trio sweep, with LAEA's 2.68×10^{-6} an order of magnitude below India's own LAEA value despite Pakistan's considerably smaller gazetteer, a point already made above. The trio-completion gap to fourth-place Sinusoidal (8.63×10^{-5}) is $4.88\times$, confirming the approximate figure already given in this chapter's interpretation. One feature not previously reported: Robinson, fifth at 1.06×10^{-4}, outperforms Equirectangular (sixth, 1.36×10^{-4}) here, the reverse of their relative order in India and in most other countries examined in this volume, where Equirectangular typically places above Robinson. The bottom of the table — Aitoff, Hammer, Van der Grinten, Mercator, tightly clustered between 4.51×10^{-4} and 5.71×10^{-4} — gives way to Winkel Tripel and, worst of all, Mollweide at 1.06×10^{-3}, the largest value recorded for this region.

Table 6.3: Pakistan: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 150,394 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	2.68×10^{-6}
2	Albers Equal-Area Conic	8.12×10^{-6}
3	Lambert Conformal Conic	1.77×10^{-5}
4	Sinusoidal	8.63×10^{-5}
5	Robinson	1.06×10^{-4}
6	Equirectangular	1.36×10^{-4}
7	Aitoff	4.51×10^{-4}
8	Hammer	5.26×10^{-4}
9	Van der Grinten	5.60×10^{-4}
10	Mercator	5.71×10^{-4}
11	Winkel Tripel	9.49×10^{-4}
12	Mollweide	1.06×10^{-3}

6.4 Interpretation

India presents the most demanding single test of the top-40-plus-seeded-40 sampling scheme in the whole study: with well over half a million populated-place candidates, the seeded-random half of the sample draws from an enormous pool, and the country's own territory spans a wide range of latitude, among the widest of any single country examined. First place is unambiguous (Table 6.2): LAEA at 1.73×10^{-5}, well clear of the rest of the field. But India is one of the three exceptions, alongside Spain and Algeria, where the regionally-fitted trio does not sweep the top three outright: Sinusoidal (1.10×10^{-4}) edges out Albers (1.14×10^{-4}) for second place, pushing LCC (1.43×10^{-4}) to fourth. Sinusoidal is, in fact, the only whole-Earth (non-regionally-fitted) projection to break into any country's top two anywhere in the entire study, plausibly because its true-scale meridians and parallels partially compensate for a strongly north–south-elongated territory even without an explicit centring parameter beyond the shared longitude shift of Section 2.4. This is the same mechanism tentatively proposed for Algeria's exception in Chapter 4 — a coastal-or-elongated settlement geometry giving Sinusoidal's true-scale-along-parallels property an unusual advantage — and India's case, given its exceptionally wide latitudinal span, is if anything a cleaner illustration of that mechanism than Algeria's own.

Pakistan, more compact in latitudinal extent than India but still substantial, shows the cleaner version of the region's pattern (Table 6.3): LAEA, Albers, and LCC occupy the first three ranks with no whole-Earth projection intruding, and the gap between LCC (third, 1.77×10^{-5}) and the best whole-Earth projection, Sinusoidal (fourth, 8.63×10^{-5}), is a factor of roughly 4.9 — smaller than the typical regional gap seen elsewhere in this volume, though the ordering itself is undisturbed. That Pakistan's fourth-place gap is comparatively narrow while India's Sinusoidal actually overtakes a member of the trio outright is consistent with a single underlying story: the more elongated or latitudinally extensive a country's territory, the more competitive Sinusoidal becomes relative to the regionally-fitted trio, with India (the more elongated of the two) crossing the threshold into an outright top-two finish and Pakistan (the more compact of the two) merely narrowing the gap without closing it. Pakistan's own LAEA value, 2.68×10^{-6}, is in fact smaller than India's by an order of magnitude despite Pakistan's considerably smaller gazetteer, once again illustrating the point already made in Chapter 4 regarding

Morocco: gazetteer size and resulting IDI magnitude are not obviously coupled once the fixed $n = 80$ sample has been drawn.

Methodological Note

The India and Pakistan results should not be over-read as evidence about sub-continental-scale mapping generally: both countries are analysed here at their full national extent with a single centroid and a single pair of standard parallels, which is exactly the regime in which a single azimuthal or conic projection is expected to perform best. A finer within-country regional breakdown (comparable to treating northern and southern India separately) is a natural further extension not attempted in this volume, and would be needed before any claim about, say, projection choice for a state-level or provincial map of either country could be supported by this pipeline.

6.5 Summary

Despite being the two most data-rich countries in the study by a wide margin, and despite India's unusually large latitudinal span, South Asia reproduces the finding already seen in every other region at the level that matters most: LAEA wins outright in both countries, by one to three orders of magnitude over the best whole-Earth alternative. India is, however, one of only three countries in the whole study (with Spain and Algeria) where the full regionally-fitted trio does not occupy the top three ranks — here because Sinusoidal, unusually well suited to India's elongated shape, displaces LCC to fourth place, and does so more decisively (an outright top-two finish) than in either of the study's two other trio-sweep exceptions. Read alongside Chapter 4's Algeria and Chapter 3's Spain, South Asia's India completes a small but consistent pattern: every one of the three countries where the trio-sweep fails does so because Sinusoidal or Winkel Tripel exploits a specific shape or settlement characteristic (elongation, coastal concentration) that a fixed pair of latitude-quartile standard parallels does not fully capture, never because a projection with no shape-adaptive property at all intrudes unexpectedly.

Chapter 7

Atlantic and Western Europe: Austria, Switzerland, the United Kingdom, Ireland

7.1 Regional character

As explained in Chapter 1, this fifth group is included for a stated reason of data availability rather than thematic continuity with the Mediterranean, Maghreb, Caucasus, or South Asia groupings: Austria and Switzerland are landlocked, mountainous Central European states; the United Kingdom and Ireland are Atlantic island nations. What they share is higher latitude ($47°$–$59°$N) than any of the other four regions, making this group a useful check on whether the central finding is latitude-dependent. It is worth being explicit about why this check matters: Mercator's distortion is a textbook function of latitude, growing without bound toward the poles, so if the volume's whole-Earth projections were failing merely because the twenty countries examined so far happened to sit at comparatively low latitudes, this region — reaching as far north as the United Kingdom's own sample extent — would be exactly where that failure mode should show through most clearly, either by narrowing the gap between the regionally-fitted trio and the whole-Earth projections, or by producing an unusually large absolute worst-case IDI value somewhere in the region.

7.2 Regional summary

7.3 Country-by-country detail

As in the preceding regional chapters, this section reports the full ranking and, where a finite one exists, the trio-completion gap for each of the four countries in this region.

Austria (Table 7.2) is the volume's most extreme resolution-floor case: LAEA, LCC, and Albers all register exactly 0, so no finite trio-completion gap can be computed here — the honest statement is that the gap is qualitative rather than numerical, separating three computational zeros from Sinusoidal's fourth-place 1.18×10^{-4}. That fourth place is itself unusual among whole-Earth projections in this region: in absolute terms it is closer to the trio than any other whole-Earth projection manages in any other Atlantic country. The remainder of Austria's table — Equirectangular, Mercator, Winkel Tripel, Van der Grinten, Mollweide, Aitoff — descends to Hammer (eleventh, 6.90×10^{-3}) and, worst of all, Robinson at 7.72×10^{-3}, second only to the United Kingdom's Hammer value as the largest worst-case figure recorded anywhere in this region.

Switzerland (Table 7.3) gives the region's narrowest non-zero trio-completion gap, $2.58\times$, from LCC (2.36×10^{-5}) to fourth-place Sinusoidal (6.08×10^{-5}) — LAEA is again an exact zero.

Table 7.1: Atlantic and Western Europe: regional mean Informational Distortion Index by projection, unweighted arithmetic mean over the 4 constituent countries (Austria, Switzerland, United Kingdom, Ireland).

Rank	Projection	Mean IDI (bits)
1	Lambert Azimuthal Equal-Area	0
2	Albers Equal-Area Conic	2.55×10^{-5}
3	Lambert Conformal Conic	7.60×10^{-5}
4	Sinusoidal	1.48×10^{-4}
5	Equirectangular	1.95×10^{-4}
6	Winkel Tripel	1.48×10^{-3}
7	Mercator	1.66×10^{-3}
8	Van der Grinten	1.71×10^{-3}
9	Aitoff	3.30×10^{-3}
10	Mollweide	3.98×10^{-3}
11	Robinson	5.56×10^{-3}
12	Hammer	5.60×10^{-3}

Albers and LCC sit within a factor of 1.03 of each other, the tightest such pairing recorded in this region. The middle table — Equirectangular, Winkel Tripel, Mercator — descends fairly evenly before Van der Grinten and Mollweide, with Aitoff, Robinson, and Hammer closing the ranking at 1.29×10^{-3}, 1.69×10^{-3}, and 1.94×10^{-3} respectively — the smallest set of worst-case values recorded for any Atlantic country.

The United Kingdom (Table 7.4) gives the region's narrowest trio-completion gap in relative terms, $1.21\times$, from LCC (2.57×10^{-4}) to fourth-place Equirectangular (3.11×10^{-4}) — itself barely ahead of fifth-place Sinusoidal at 3.18×10^{-4}, a near-tie for fourth not seen in any other Atlantic country. This is also, as already noted, the country whose exact-zero LAEA result is hardest to attribute to compactness alone, given the United Kingdom's comparatively large and well-documented gazetteer. The bottom half of the table — Winkel Tripel, Van der Grinten, Aitoff, Mercator, Mollweide, Robinson — descends steadily to Hammer's 9.02×10^{-3}, the single largest worst-case value recorded anywhere in the Atlantic region and, as already noted, still smaller than the two Mediterranean values that hold the volume-wide record.

Ireland (Table 7.5) shares Austria's and the United Kingdom's exact-zero LAEA result and adds a second zero of its own, Albers, leaving LCC (2.35×10^{-5}) as the only non-zero trio member. The trio-completion gap, from LCC to fourth-place Sinusoidal (9.41×10^{-5}), is $4.00\times$. The remaining table — Equirectangular, Mercator, Van der Grinten, Winkel Tripel — descends to a tightly packed bottom four, Aitoff, Mollweide, Hammer, and Robinson, all within a comparatively narrow band between 4.09×10^{-3} and 4.71×10^{-3}, with Robinson worst at 4.71×10^{-3} — consistent with this chapter's observation that Robinson and Hammer trade the region's worst two positions without following the landlocked/island geographic split.

7.4 Interpretation

The regional mean (Table 7.1) reproduces the by-now-familiar ordering, with the most extensive cluster of exact-zero IDI values seen anywhere in the study: LAEA returns exactly 0 in *all four* countries of this region (Tables 7.2 to 7.5), the only region where this happens for every one of its constituent countries. Austria goes further still, with LCC and Albers also returning exactly 0 alongside LAEA — three simultaneous zeros, the largest number recorded for any

Table 7.2: Austria: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 20,913 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	0
2	Lambert Conformal Conic	0
3	Albers Equal-Area Conic	0
4	Sinusoidal	1.18×10^{-4}
5	Equirectangular	1.47×10^{-4}
6	Mercator	2.15×10^{-4}
7	Winkel Tripel	1.13×10^{-3}
8	Van der Grinten	2.42×10^{-3}
9	Mollweide	3.22×10^{-3}
10	Aitoff	4.44×10^{-3}
11	Hammer	6.90×10^{-3}
12	Robinson	7.72×10^{-3}

single country in the study, and a considerably more extreme instance of the resolution-floor phenomenon already noted for Armenia and Georgia in Chapter 5. Ireland returns two exact zeros (LAEA and Albers); Switzerland and the United Kingdom return one each (LAEA only). This cluster is consistent with Austria and Switzerland being among the most territorially compact countries in the twenty-one-country study, and reinforces the reading, already given in Chapter 5, that these are a resolution artefact of country compactness interacting with a fixed 25-bin histogram, not a claim that any projection is literally, continuously distortion-free — a reading that the United Kingdom's exact zero, despite being far larger and less compact than Austria or Switzerland, shows cannot be reduced to compactness alone. The United Kingdom's own gazetteer, at 43,712 feature-class-P records, is in fact the second-largest of any country in this region and larger than every Mediterranean European country except France and Italy, which makes its exact-zero LAEA result the volume's clearest illustration that the resolution floor is a function of the fitted projection's own residual distortion at this specific sample size and bin count, not simply a function of how small or densely documented the country happens to be.

Hammer and Robinson trade the worst and second-worst positions throughout this region, but the split does not follow the landlocked/island geographic line one might expect: Hammer is worst for Switzerland and the United Kingdom, while Robinson is worst for Austria and Ireland, with the other of the pair always taking second-worst in each case — Austria (landlocked) and Ireland (an Atlantic island) share a worst projection, as do Switzerland (landlocked) and the United Kingdom (an Atlantic island), cutting directly across the geographic grouping rather than following it. This is a small but genuine piece of evidence against reading too much geographic determinism into which whole-Earth projection loses worst in a given country: whatever drives Hammer's or Robinson's relative failure here tracks something more specific to each country's own shape and latitude than the landlocked-versus-island distinction that otherwise organizes this chapter's country roster. Nor is this the region with the highest worst-case IDI values in the study: the United Kingdom's Hammer value of 9.02×10^{-3} is the largest recorded in this region, but it is smaller than two values already reported in Chapter 3 — Spain's Van der Grinten value of 2.08×10^{-2} and Portugal's Hammer value of 1.85×10^{-2} — so no simple claim that whole-Earth projections degrade monotonically with latitude is supported by this dataset. If anything, the comparison runs the other way: two mid-latitude

Table 7.3: Switzerland: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 13,199 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	0
2	Albers Equal-Area Conic	2.30×10^{-5}
3	Lambert Conformal Conic	2.36×10^{-5}
4	Sinusoidal	6.08×10^{-5}
5	Equirectangular	1.26×10^{-4}
6	Winkel Tripel	1.67×10^{-4}
7	Mercator	2.08×10^{-4}
8	Van der Grinten	4.87×10^{-4}
9	Mollweide	7.29×10^{-4}
10	Aitoff	1.29×10^{-3}
11	Robinson	1.69×10^{-3}
12	Hammer	1.94×10^{-3}

Mediterranean countries produce worse whole-Earth-projection outcomes than any country in this considerably higher-latitude region, which is the opposite of what a naive latitude-driven-failure story would predict, and is consistent instead with the volume's broader finding that a country's specific shape and settlement pattern, not its latitude band as such, is what principally governs how badly a whole-Earth compromise projection performs against it.

7.5 Summary

Even in a region assembled for reasons of data convenience rather than geographic coherence, and at the highest latitudes tested anywhere in this study, the central finding holds without exception at first place: LAEA wins outright in all four countries, dominating every whole-Earth compromise projection tested by one to several orders of magnitude. The region's worst-projection pattern, however, does not sort neatly by geography, and its highest worst-case values, while large, are not the largest recorded in the volume — two Mediterranean countries hold that distinction instead. Taken together, these two observations answer the question this chapter set out to ask: the volume's central regularity is not an artefact of the mostly mid-latitude sample examined in Chapters 3 to 6, since it survives, unchanged in its first-place ordering, at the highest latitudes tested anywhere in the study, even as the finer details of which whole-Earth projection loses worst continue to depend on country-specific shape rather than on latitude alone.

Table 7.4: United Kingdom: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 43,712 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	0
2	Albers Equal-Area Conic	7.88×10^{-5}
3	Lambert Conformal Conic	2.57×10^{-4}
4	Equirectangular	3.11×10^{-4}
5	Sinusoidal	3.18×10^{-4}
6	Winkel Tripel	1.87×10^{-3}
7	Van der Grinten	2.15×10^{-3}
8	Aitoff	3.36×10^{-3}
9	Mercator	5.74×10^{-3}
10	Mollweide	7.78×10^{-3}
11	Robinson	8.14×10^{-3}
12	Hammer	9.02×10^{-3}

Table 7.5: Ireland: Informational Distortion Index by projection, $n = 80$ sampled populated places, 3160 pairwise distances (out of 12,159 feature-class-P records nationwide). Lower is better.

Rank	Projection	IDI (bits)
1	Lambert Azimuthal Equal-Area	0
2	Albers Equal-Area Conic	0
3	Lambert Conformal Conic	2.35×10^{-5}
4	Sinusoidal	9.41×10^{-5}
5	Equirectangular	1.96×10^{-4}
6	Mercator	4.80×10^{-4}
7	Van der Grinten	1.77×10^{-3}
8	Winkel Tripel	2.74×10^{-3}
9	Aitoff	4.09×10^{-3}
10	Mollweide	4.20×10^{-3}
11	Hammer	4.54×10^{-3}
12	Robinson	4.71×10^{-3}

Chapter 8

Synthesis: One Robust Regularity Across Twenty-One Countries

8.1 The result, stated once, plainly

Across five regions chosen for geographic, climatic, and cultural diversity — Mediterranean Europe, the Maghreb, the Caucasus and Anatolia, South Asia, and (for reasons of data availability rather than thematic fit) Atlantic and Western Europe — and across twenty-one countries individually, the same ordering recurs with no exception: a projection centred and fitted to the region using only information available at mapmaking time (its own centroid, and for conic projections its own latitude quartiles) achieves an Informational Distortion Index one to three orders of magnitude lower than any projection designed for whole-Earth compromise display. Table 8.1 gives the grand mean across all twenty-one countries, unweighted by region size.

Table 8.1: Grand mean Informational Distortion Index across all twenty-one countries, unweighted arithmetic mean of the 21 country-level IDI values for each projection (not a repooling of raw points, and not weighted by region size).

Rank	Projection	Grand mean IDI (bits)
1	Lambert Azimuthal Equal-Area	1.52×10^{-5}
2	Albers Equal-Area Conic	4.34×10^{-5}
3	Lambert Conformal Conic	1.54×10^{-4}
4	Sinusoidal	6.68×10^{-4}
5	Equirectangular	8.19×10^{-4}
6	Winkel Tripel	1.81×10^{-3}
7	Mollweide	2.29×10^{-3}
8	Robinson	2.65×10^{-3}
9	Aitoff	2.69×10^{-3}
10	Van der Grinten	2.81×10^{-3}
11	Mercator	2.82×10^{-3}
12	Hammer	4.15×10^{-3}

The ratio between the best grand-mean projection (LAEA, 1.52×10^{-5}) and the worst (Hammer, 4.15×10^{-3}) is a factor of 274. *First* place among the twenty-one countries never goes to a whole-Earth compromise projection: LAEA, Albers, or LCC is first in every single one. LAEA takes first place in nineteen of the twenty-one; the only two exceptions are Turkey, where LCC edges it out by the narrowest margin recorded anywhere in the study (a ratio of 1.019), and Portugal, where Albers wins outright. The stronger claim, that the full regionally-fitted trio always occupies the top three ranks, is not universal: it fails in three countries — Spain, Algeria, and India — where Winkel Tripel or Sinusoidal intrudes into third place, each time displacing whichever member of the trio is weakest for that country's specific shape. Chapters 3, 4 and 6

report each case individually; this chapter's point is only that the exceptions are few, specific, and never touch first place.

8.2 Cross-regional rank stability

Table 8.2 lists, for every projection, its rank (1=lowest IDI, 12=highest) by regional mean in each of the five regions. The top three ranks (LAEA, Albers, LCC, in some order) are essentially invariant across regions; Hammer occupies the bottom rank in three of the five regions (Maghreb, Caucasus and Anatolia, Atlantic and Western Europe) and is never better than eleventh in the other two. The middle of the table — Sinusoidal, Equirectangular, Mollweide, Winkel Tripel, Mercator, Van der Grinten, Aitoff, Robinson — reshuffles considerably more from region to region, which is itself informative: whole-Earth compromise projections do not all fail for the same reason or in the same relative order once population-weighted real data replaces a synthetic grid, whereas regionally-fitted projections succeed for essentially the same reason (local recentring) everywhere.

Table 8.2: Cross-regional rank of each projection by regional mean IDI (1 = lowest IDI, 12 = highest), one column per region, ordered by overall mean rank.

Projection	Med. Eur.	Maghreb	Cauc. & Anat.	South Asia	Atlantic/W. Eur.
Lambert Azimuthal Equal-Area	2	1	1	1	1
Albers Equal-Area Conic	1	2	3	2	2
Lambert Conformal Conic	3	3	2	3	3
Sinusoidal	5	4	5	4	4
Equirectangular	7	5	6	5	5
Winkel Tripel	4	8	7	11	6
Mollweide	6	7	4	12	10
Aitoff	8	10	11	6	9
Robinson	9	6	10	8	11
Mercator	12	9	8	10	7
Van der Grinten	10	11	9	9	8
Hammer	11	12	12	7	12

8.3 Regional variation in the grand-mean effect size

The $274\times$ grand-mean ratio quoted above is a single number computed across all twenty-one countries at once, and it is worth asking whether that figure is representative of each region individually or is instead an artefact of averaging across regions whose own effect sizes differ considerably. Table 8.3 answers this directly, computing the same best-to-worst regional-mean ratio separately within each of the five regions.

Four of the five regions cluster in a comparatively narrow band, from South Asia's $175.7\times$ to Maghreb's $237.6\times$, all reasonably close to, and generally somewhat below, the $274\times$ grand-mean figure. The Caucasus and Anatolia is a clear outlier at $657.9\times$, more than double the next-widest regional ratio and more than triple the narrowest. This is driven specifically by the numerator side of the ratio rather than an unusually poor worst-case performer: the region's best-projection mean, 2.969×10^{-6}, is pulled down by Armenia's and Georgia's own near-zero LAEA and Albers values (Chapter 5), while its worst-projection mean, Hammer at $1.953\ \times$

Table 8.3: Best/worst projection ratio at the regional-mean level, for each of the five regions, alongside the twenty-one-country grand mean already given in Table 8.1. Atlantic and Western Europe's best-projection column is LAEA at an exact regional-mean 0 (all four constituent countries individually return exact-zero LAEA, Chapter 7); the ratio shown for that region instead uses Albers, the best *non-zero* regional-mean projection, since a ratio against an exact zero is undefined.

Region	Best (mean IDI)	Worst (mean IDI)	Ratio
Mediterranean Europe	Albers (2.759×10^{-5})	Mercator (5.441×10^{-3})	197.2×
Maghreb	LAEA (1.994×10^{-5})	Hammer (4.737×10^{-3})	237.6×
Caucasus and Anatolia	LAEA (2.969×10^{-6})	Hammer (1.953×10^{-3})	657.9×
South Asia	LAEA (1.001×10^{-5})	Mollweide (1.758×10^{-3})	175.7×
Atlantic and Western Europe	Albers† (2.545×10^{-5})	Hammer (5.600×10^{-3})	220.0×
All twenty-one (grand mean)	LAEA (1.52×10^{-5})	Hammer (4.15×10^{-3})	274.1×

10^{-3}, is unremarkable and in fact smaller than three of the other four regions' own worst-projection means. The 274× grand-mean figure is therefore not a typical regional outcome so much as a number pulled upward by one region's unusually small denominator; a reader who wants a sense of the "typical" regional effect size is better served by the 175–240× range that four of the five regions actually occupy than by the grand mean alone.

Methodological Note

The Atlantic and Western Europe row requires an honest caveat already flagged in Chapter 7: its true best-projection regional mean, LAEA, is an exact 0, since all four constituent countries individually return exact-zero LAEA at this study's 25-bin resolution. A ratio against an exact zero is undefined, not merely large, so the 220.0× figure reported for this region uses Albers, the best *non-zero* regional-mean projection, as a practical substitute. This means the Atlantic figure in Table 8.3 is not directly comparable to the other four rows on a like-for-like basis — it understates, rather than overstates, how favourable LAEA's true performance is in this region.

8.4 The trio's internal ordering, country by country

Chapters 3 to 7 each report, for their own countries, which of LAEA, Albers, and LCC takes first place; this section assembles that same internal ordering — first, second, *and* third among the trio itself — across all twenty-one countries at once, in Table 8.4, and asks whether the runner-up and third-place positions follow any pattern of their own.

The pattern is considerably more regular than the first-place tally alone suggests. In fifteen of the twenty-one countries, the trio's internal order is identical — LAEA first, Albers second, LCC third — and this is not a geographically random fifteen: it comprises five of the six Mediterranean countries (all but Portugal), four of the five Maghreb countries (all but Algeria), only one of the four Caucasus and Anatolia countries (Armenia), both South Asian countries, and three of the four Atlantic countries (all but Austria). The remaining six countries — Portugal, Algeria, Azerbaijan, Georgia, Turkey, and Austria — all share a single structural feature that the "typical" fifteen do not: in each of them, LCC outranks Albers, either narrowly (Algeria, Azerbaijan, Georgia, Austria, where LCC takes second place ahead of Albers) or outright (Portugal, Turkey, where LCC takes first place). Three of these six are Caucasus and

Table 8.4: The regionally-fitted trio's internal first/second/third ordering in each of the twenty-one countries. The three sweep-failure countries (Spain, Algeria, India, Chapters 3, 4 and 6) are marked; their trio-internal order is still well-defined even though a whole-Earth projection intrudes into what would otherwise be the trio's own third rank. The final row's three entries in each column give the LAEA/Albers/LCC split of that rank across all twenty-one countries.

Country	1st	2nd	3rd
Italy	LAEA	Albers	LCC
Spain*	LAEA	Albers	LCC
Portugal	Albers	LCC	LAEA
Croatia	LAEA	Albers	LCC
Slovenia	LAEA	Albers	LCC
France	LAEA	Albers	LCC
Algeria*	LAEA	LCC	Albers
Morocco	LAEA	Albers	LCC
Tunisia	LAEA	Albers	LCC
Libya	LAEA	Albers	LCC
Mauritania	LAEA	Albers	LCC
Armenia	LAEA	Albers	LCC
Azerbaijan	LAEA	LCC	Albers
Georgia	LAEA	LCC	Albers
Turkey	LCC	LAEA	Albers
India*	LAEA	Albers	LCC
Pakistan	LAEA	Albers	LCC
Austria	LAEA	LCC	Albers
Switzerland	LAEA	Albers	LCC
United Kingdom	LAEA	Albers	LCC
Ireland	LAEA	Albers	LCC
Count	19/1/1	1/15/5	1/5/15

Anatolia countries — the only region in the volume where LCC displaces Albers from second place in a majority of its constituent countries — and Portugal and Austria are, respectively, the region's most elongated Mediterranean territory and an Atlantic country whose trio ordering is dominated by a resolution-floor effect rather than shape, two countries that otherwise have little geographically in common with the Caucasus.

Methodological Note

This six-country cluster should not be read as evidence of a single unified mechanism: Portugal's own account (Chapter 3) attributes LCC's strength there to elongation, exactly the mechanism tested formally, and not confirmed at conventional significance, in Section 9.4; Turkey's and the Caucasus cluster's account (Chapter 5) instead emphasises the region's mountainous mid-latitude character generally, without singling out elongation specifically; and Austria's case is dominated by a resolution-floor effect (Chapter 7) in which LAEA, Albers, and LCC are all so close to a computational zero that their relative order may owe as much to which bin a handful of borderline pairwise distances happen to fall into as to any geometric property of the country itself. What the six countries share is only the qualitative outcome — LCC over Albers — not a demonstrated common cause; Table 8.4's contribution is to make that outcome, and its minority status (six of twenty-one), precise and countable, not to explain it definitively.

8.5 What drives the effect, and what does not

The mechanism behind this regularity is not mysterious and does not require the fragile second-order analytic argument attempted in Chapter 10: LAEA, Albers, and LCC are, by construction, given the sample's own centroid (and, for the two conic projections, the sample's own latitude quartiles as standard parallels) before a single distance is computed, while Equirectangular, Mercator, Sinusoidal, Robinson, and Winkel Tripel either have no free centring parameter at all in their published definitions or are evaluated at a shared, region-independent standard parallel (Winkel Tripel's $\varphi_1 = \arccos(2/\pi)$ is fixed by definition, not fitted). This is closer to a statement about *which projections in this study are given country-specific information and which are not* than a statement about which projection family is intrinsically superior in some projection-independent sense — a distinction already anticipated in Section 2.4 and worth restating here plainly, since it is the single most important caveat to the whole volume's central finding.

Methodological Note

This is not a criticism of the finding, only a precise statement of what it does and does not establish. It establishes: *given* the freedom to recentre and refit standard parallels to a specific country, the equal-area azimuthal and conic families recover that freedom's benefit almost entirely, while the whole-Earth pseudocylindrical and compromise families in this study do not expose an equivalent degree of freedom to exploit, even where they nominally could (Aitoff, Mollweide, Hammer, and Van der Grinten all take a λ_0 parameter here and were recentred exactly as the others were, per Section 2.4, yet still trail badly). It does *not* establish that no whole-Earth-style projection formula could ever be adapted to match LAEA or Albers at national scale — only that the twelve standard, published formulas tested here, evaluated as published (with recentring where the published formula allows it), do not.

8.6 Relation to the failed scaling-law attempt

Chapter 10 documents, at full length and without euphemism, why the more ambitious analytic route — an asymptotic scaling law for the IDI as a sampled region shrinks toward a point — was abandoned. It is worth stating explicitly why that failure does not undermine the finding of this chapter. The scaling-law attempt concerned the *local*, infinitesimal-neighbourhood behaviour of the IDI at a single point, governed by second-order metric expansions and vulnerable to a histogram-estimator noise floor precisely because the signal it was chasing shrinks to nothing as $\eta \to 0$. The finding of this chapter concerns the IDI at full national scale, where geodesic distances range from a few kilometres to (for India) several thousand, giving the JSD estimator a wide, well-populated distance range to resolve — the opposite regime from the one in which the scaling-law estimator broke down. The two studies are not in tension; they simply operate at scales where different failure modes are, and are not, active.

8.7 Summary and closing remark

This volume set out to answer a single, deliberately modest question left open by the parent monograph: computed on real settlement data rather than synthetic grids, does the Informational Distortion Index tell a coherent story? It does, and the story is simple enough to

state in one sentence: at national and regional scale, across twenty-one countries spanning five culturally and geographically distinct regions, a projection fitted to the mapped region wins first place in every single country, typically by one to two orders of magnitude over the best whole-Earth alternative and by a factor of 274 at the level of the twenty-one-country grand mean, with the only genuine competition occurring between the fitted projections themselves — and, in three specific and individually reported countries, with Sinusoidal or Winkel Tripel for third place. Reaching this result required abandoning, honestly and in print, a more ambitious theoretical claim that did not survive its own numerical test — a discipline the parent volume already modelled in its treatment of the Aitoff/Hammer ranking instability and the non-significant Tissot correlation, and one this companion volume has tried to extend rather than relax.

Chapter 9

Robustness Diagnostics

Every result in Chapters 3 to 8 rests on two fixed methodological choices stated in Chapter 2: a 25-bin histogram, and an $n = 80$ sample built as top-40-by-population plus a seeded-random 40 more. Neither number was tuned to produce the reported ranking; both were fixed before any projection was evaluated, and both could, in principle, have been chosen differently without any change to the pipeline's logic. That is precisely what makes them worth interrogating rather than simply asserting. A finding that survives only at one particular bin count, or only at one particular sample size, is a much weaker claim than one that survives across a range of reasonable choices, and a monograph that reports the former as if it were the latter would be quietly overselling its own result. This chapter exists to close that gap: it does not introduce new data, new countries, or new projections, but it does introduce five genuinely new computations, run once each on the same real GeoNames coordinates and pipeline output used throughout the volume, with no numerical entry adjusted by hand after the fact.

The five checks are logically independent of one another and are aimed at five different soft spots in the argument built up over the preceding chapters. The first two ask the most direct question a skeptical reader could pose: if the 25-bin histogram or the $n = 80$ sample had been chosen differently, would the headline result — that a regionally-fitted projection takes first place in every one of the twenty-one countries — have survived? The third asks whether the qualitative impression of cross-regional agreement built up narratively across Chapters 3 to 7 and summarised in Chapter 8 holds up once it is reduced to a single, pre-registered statistic rather than left as an interpretive gloss on five separate rank tables. The fourth is different in character: it does not defend a finding already made but tests a hypothesis that was only ever advanced impressionistically, that a country's elongation predicts how closely the conic Lambert Conformal Conic projection can approach the azimuthal Lambert Azimuthal Equal-Area. As will become clear in Section 9.4, this last check does not confirm the hypothesis, and the chapter reports that outcome as plainly as it reports the three checks that do go the volume's way. The fifth, taken up in Section 9.5, formalises a different impressionistic claim repeated across several regional chapters — that gazetteer size does not predict the trio-completion gap introduced in Chapters 3 to 7's country-by-country detail sections — and finds, this time, a clean null result consistent with what those chapters had each suggested in passing. The full script implementing all five checks is reproduced in Section 11.4, so that every number quoted below can be traced back to its source rather than taken on faith.

9.1 Bin-count sensitivity

The 25-bin histogram used throughout this volume is inherited directly from the parent monograph's own convention, and like that convention it is a fixed, round-number choice rather than a value selected because it happened to produce a clean ranking. Histogram-based Jensen–Shannon divergence estimators are known to be sensitive to bin count in exactly the regime this volume operates in — moderate sample sizes ($n = 80$ pairwise distances collapse to $\binom{80}{2} = 3160$ pairs, which is generous, but not enormous, once spread across 25 or more bins)

— so the natural first question is whether the 25-bin choice is doing any of the argumentative work. Table 9.1 answers this by rerunning the entire twenty-one-country pipeline, unchanged in every other respect (same coordinates, same seeded-random component, same centring and recentring logic from Chapter 2), at three additional bin counts chosen to bracket the baseline from both directions: 15 (visibly coarser), 35, and 50 (visibly finer, and close to the practical limit at which individual bins begin to contain very few pairs for the smaller countries in the sample).

Table 9.1: Bin-count sensitivity: grand-mean top-3 projections and first-place stability, at $N_{\text{bins}} \in \{15, 25, 35, 50\}$ (25 is the baseline used throughout Chapters 3 to 8).

N_{bins}	Grand-mean top 3 (order)	First-place flips vs. baseline	Trio-sweep exceptions
15	LAEA, Albers, LCC	3 (all within trio)	CH, DZ, GB, GE, IN
25 (baseline)	LAEA, Albers, LCC	—	ES, DZ, IN
35	LAEA, Albers, LCC	1 (within trio)	DZ, GB, IN, LY
50	LAEA, Albers, LCC	1 (within trio)	DZ, GB, IN

The headline finding turns out to be remarkably insensitive to this choice. At all four bin counts tested, the grand-mean ranking of the top three projections is identical in both membership and order — LAEA first, Albers second, LCC third — and, more importantly for the volume's central claim, no whole-Earth compromise projection ever climbs into first place in any single country at any bin count tested. Coarsening the histogram to 15 bins does not, for instance, let Robinson or Winkel Tripel sneak past a regionally-fitted projection anywhere; it only reshuffles the internal ordering of the trio itself, and even that reshuffling is limited. Table 9.2 makes this precise by listing every individual country in which the identity of the first-place projection changes as the bin count moves away from 25: there are five such changes across three alternative resolutions combined (15, 35, and 50), and in every single one of them the reassignment is a swap between two members of the LAEA/Albers/LCC trio, never a promotion of a projection from outside it.

Table 9.2: Every individual first-place reassignment observed across the three alternative bin counts, relative to the $N_{\text{bins}} = 25$ baseline. In every case the reassignment is between two members of the regionally-fitted trio; no whole-Earth projection takes first place at any tested bin count, in any country.

N_{bins}	Country	Baseline first	New first
15	Portugal	Albers	LAEA
15	Libya	LAEA	Albers
15	United Kingdom	LAEA	LCC
35	Austria	LAEA	LCC
50	Spain	LAEA	Albers

The secondary finding of the volume — the narrower and more fragile claim that the full regionally-fitted trio sweeps the top three ranks, rather than merely occupying first place, in eighteen of the twenty-one countries — turns out to be visibly less robust to this same manipulation, and it is worth being explicit about why that asymmetry is not itself a weakness in the argument. At the 25-bin baseline, the trio-sweep fails in exactly three countries, Spain, Algeria, and India, each of which is discussed individually in Chapters 3, 4 and 6. Rerunning at 15, 35, and 50 bins does not reproduce this same exception set: at 15 bins the exceptions

become Switzerland, Algeria, the United Kingdom, Georgia, and India; at 35 bins they become Algeria, the United Kingdom, India, and Libya; and at 50 bins they narrow to Algeria, the United Kingdom, and India. Pooling the three alternative-resolution exception sets and removing the two countries (Algeria and India) that already appear in the baseline exception set leaves four countries — Switzerland, the United Kingdom, Georgia, and Libya — each of which enters some alternative-resolution exception list without ever appearing in the 25-bin baseline list. This is the mechanism already diagnosed narratively in Chapters 5 and 7: the margin separating the weakest member of the regionally-fitted trio from the best-placed whole-Earth intruder is frequently small enough, in absolute IDI terms, that it can flip sign depending on exactly where the histogram bin edges happen to fall relative to the underlying pairwise-distance distribution, a sensitivity that has nothing to do with the projections' intrinsic merits and everything to do with where a discretisation boundary lands. The conclusion this chapter draws from that instability is correspondingly narrow and, it should be said, more honest than simply repeating the baseline exception list as if it were a stable geographic fact: the qualitative statement that *some* small number of countries will break the clean trio-sweep at any given resolution is robust across all four bin counts tested, but *which* countries do so is itself a function of the bin count, and should not be read as telling us anything specific about Spain's, Algeria's, or India's geography as opposed to Switzerland's, the United Kingdom's, Georgia's, or Libya's.

9.2 Sample-size and sample-composition sensitivity

The second arbitrary methodological choice inherited from Chapter 2 is the $n = 80$ sample itself, built as the top 40 settlements by population plus a seeded-random 40 drawn from the remainder. This composition was chosen to balance two competing concerns: relying on population rank alone would bias the sample toward a handful of dominant metropolitan centres and understate the geographic spread that a projection's recentring is actually meant to be tested against, while relying on a purely random sample would risk including so many minor settlements that population-weighted real-world usage patterns are lost entirely. The 40/40 split was fixed before any projection was evaluated, but fixing a compromise does not establish that the compromise was necessary; Table 9.3 tests two alternatives at the extremes of what a reasonable reader might have chosen instead, keeping the 25-bin histogram and the same deterministic per-country seed used everywhere else in the volume fixed throughout. The first alternative, $n = 40$, discards the seeded-random half entirely and keeps only the top-40-by-population subset, isolating what happens when the sample is composed exclusively of a country's largest cities. The second, $n = 160$, keeps the same top-40 core but triples the seeded-random component to 120, testing whether enriching the sample with more of the dispersed, lower-population settlements it was already drawing on strengthens, weakens, or leaves untouched the central result.

As with the bin-count check, the grand-mean ordering of the trio — LAEA, Albers, LCC, in that order — survives both alternative sampling schemes without alteration, and no whole-Earth projection reaches first place in any of the twenty-one countries under either scheme. Table 9.4 lists every individual within-trio reassignment observed, and once again every single reassignment stays inside the trio: the $n = 40$ scheme produces five such reassignments (Italy, Tunisia, Libya, Turkey, and the United Kingdom each swap their first-place projection within the trio), and the $n = 160$ scheme produces three (France, Libya, and Ireland).

The asymmetry between these two counts is itself informative and deserves a mechanistic explanation rather than a bare tally. The $n = 40$ top-only scheme, which removes the random component entirely, produces the largest number of within-trio reassignments of any variant

Table 9.3: Sample-size and sample-composition sensitivity. $n = 40$ (top-only) drops the seeded-random half entirely; $n = 160$ triples the seeded-random component (top-40 + seeded-120). Both use the same deterministic per-country seed as the $n = 80$ baseline.

Sampling scheme	Grand-mean top 3 (order)	First-place flips vs. baseline	Trio-sweep exceptions
$n = 40$ (top-only)	LAEA, Albers, LCC	5 (all within trio)	AM, FR, GE, HR, IN, SI, TN
$n = 80$ (baseline)	LAEA, Albers, LCC	—	ES, DZ, IN
$n = 160$ (top-40+seed-120)	LAEA, Albers, LCC	3 (all within trio)	DZ, GB, IN, TN

Table 9.4: Every individual first-place reassignment observed under the two alternative sampling schemes, relative to the $n = 80$ baseline. As with bin-count sensitivity (Table 9.2), every reassignment stays within the regionally-fitted trio.

Scheme	Country	Baseline first	New first
n=40 (top-only)	Italy	LAEA	Albers
n=40 (top-only)	Tunisia	LAEA	Albers
n=40 (top-only)	Libya	LAEA	Albers
n=40 (top-only)	Turkey	LCC	Albers
n=40 (top-only)	United Kingdom	LAEA	Albers
n=160 (top-40+seed-120)	France	LAEA	Albers
n=160 (top-40+seed-120)	Libya	LAEA	Albers
n=160 (top-40+seed-120)	Ireland	LAEA	Albers

tested in this chapter, and this is not a surprising or worrying result once the composition of what was removed is made explicit: dropping forty of the eighty points removes exactly the geographically dispersed, lower-population settlements that a purely azimuthal or conic recentring is supposed to be tested against in the first place, leaving behind a sample weighted almost entirely toward a handful of the largest cities, which tend by their nature to cluster more tightly in space than the fuller settlement pattern does. A sample of only the largest cities is, in effect, a lower-resolution probe of a country's true geographic extent, so it is unsurprising that the finer distinctions between the three regionally-fitted projections — distinctions that depend on exactly how the sample's centroid and latitude quartiles interact with the country's true shape — become noisier and more liable to flip under this leaner sample. The $n = 160$ scheme runs in the opposite direction: rather than removing dispersed points, it adds more of them, and the result is not merely fewer reassignments (three against five) but, if anything, a sharpening of the gap between the trio and the whole-Earth projections. The grand-mean IDI for LAEA falls from 1.52×10^{-5} at the $n = 80$ baseline to 4.76×10^{-6} at $n = 160$, a more-than-threefold reduction that indicates the regionally-fitted projection's advantage becomes, if anything, more pronounced as the sample is enriched with a fuller picture of a country's settlement geography rather than a weaker one confined to its largest cities. Both directions of this sensitivity check therefore point the same way: the central finding strengthens, or at the very least holds steady, as the sample grows or is enriched with more dispersed points, and nothing in this chapter suggests that the particular figure of $n = 80$ chosen in Chapter 2 was doing any load-bearing work that a smaller or larger sample would not also do.

9.3 Cross-regional rank concordance

Chapter 8 already observes, in narrative terms, that the five regional rank tables assembled in Table 8.2 agree closely with one another at the top and bottom of the twelve-projection ranking while reshuffling considerably more in the middle. That is a defensible reading of five tables placed side by side, but it is also exactly the kind of impressionistic claim that a numerically-minded reader is entitled to distrust until it is reduced to a single, pre-specified statistic computed the same way every time. Kendall's coefficient of concordance W exists precisely for this purpose: it treats each of the five regions as an independent "judge" ranking the same twelve projections by regional mean IDI, and asks how much those five independent rankings agree with one another, on a scale running from 0 (no agreement beyond what chance would produce) to 1 (perfect agreement across every judge on every rank). Formally,

$$W = \frac{12S}{m^2(k^3 - k)}, \qquad S = \sum_{p=1}^{k} \left(R_p - \bar{R}\right)^2, \tag{9.1}$$

where $m = 5$ is the number of regions acting as judges, $k = 12$ is the number of projections being ranked, R_p is the sum of the ranks assigned to projection p across all five regions, and $\bar{R} = m(k+1)/2$ is the mean of those rank sums under the null hypothesis of no agreement at all. Computed directly from the same regional-mean rankings that already underlie Table 8.2 in Chapter 8 — so that this statistic is not an independent re-derivation from scratch but a formalisation of the exact comparison already presented there — this gives

Key Result

$W = 0.799$ on $m = 5$ judges and $k = 12$ items.

A value of 0.799 sits well above the midpoint of the $[0, 1]$ scale and indicates strong, though not perfect, agreement among the five regions on how the twelve projections should be ordered. It is worth being precise about what this number does and does not confirm. It confirms, quantitatively rather than by eye, that the five regional rankings are not simply five unrelated orderings that happen to look similar by coincidence; a W this far from 0 would be extremely unlikely to arise if each region's ranking were an independent random permutation of the twelve projections. It does not, on the other hand, claim that the five regions agree on everything: $W = 1$ is a demanding standard requiring identical rankings from every judge, and 0.799 leaves visible room for exactly the kind of disagreement already described qualitatively in Chapter 8 — the top three ranks (some ordering of LAEA, Albers, and LCC) and the bottom rank (almost always Hammer) are essentially fixed across all five regions, while the middle eight projections, spanning Sinusoidal, Equirectangular, Mollweide, Winkel Tripel, Mercator, Van der Grinten, Aitoff, and Robinson, reorder considerably from one region to the next. That pattern of firm agreement at the extremes and looser agreement in the middle is exactly the kind of partial disagreement that would pull a concordance statistic below 1 without pushing it anywhere near 0, and 0.799 is consistent with that description almost by construction.

9.4 Country shape and the LCC/LAEA margin

Chapters 4 and 5 both suggest, in passing and without any formal statistical support, that a country's elongation — a long, thin territory rather than a compact, roughly disc-shaped one — favours the conic Lambert Conformal Conic projection over the azimuthal Lambert Azimuthal Equal-Area. The intuition behind that suggestion is geometrically plausible: a conic projection

is defined by two standard parallels that can be placed to correct distortion preferentially along one preferred axis, which is exactly the kind of correction a long, thin country would benefit from, whereas an azimuthal projection corrects distortion symmetrically around a single centre point and has no comparable way to favour one axis over another. Two or three anecdotal cases — Turkey's narrow LCC win in Chapter 5 chief among them — are consistent with this story, but two or three cases do not constitute evidence of a general relationship, and the honest thing to do with an intuition this specific is to test it against the full twenty-one-country dataset rather than let it stand on the strength of its most favourable examples.

For each of the twenty-one countries, a longitude/latitude bounding-box elongation ratio is computed directly from the same $n = 80$ sample used throughout this volume, $\text{elong} = (\lambda_{\max} - \lambda_{\min})/(\varphi_{\max} - \varphi_{\min})$, giving a single number that grows with how much wider a country's longitudinal spread is relative to its latitudinal spread. This is correlated, via Spearman's rank correlation — chosen over Pearson's because the relationship being tested is directional and monotonic rather than assumed to be linear, and because the underlying IDI values span several orders of magnitude and are not well suited to a linear-correlation assumption — against $\ln(\text{IDI}_{\text{LCC}}/\text{IDI}_{\text{LAEA}})$, the log-ratio by which LCC under- or over-performs LAEA for that country, where negative values mean LCC is closer to, or has actually overtaken, LAEA.

Methodological Note

Nine of the twenty-one countries — Croatia, Slovenia, Tunisia, Armenia, Georgia, Austria, Switzerland, the United Kingdom, and Ireland — return an exact IDI of 0 for LAEA or LCC or both at the $n = 80$, 25-bin baseline, a resolution-floor artefact already discussed in Chapters 5 and 7 and revisited from a different angle in Section 9.1 above. An exact zero in either numerator or denominator makes the log-ratio undefined, so these nine countries are excluded from the correlation that follows, leaving $n = 12$ valid countries. This is a real limitation of what the available data can support at this histogram resolution, disclosed here rather than papered over, and not a filtering choice made after the fact to improve the correlation's apparent strength.

Key Result

Spearman $\rho = -0.245$ ($p = 0.443$, $n = 12$).

The correlation that emerges is weak, points in the direction the narrative discussion in Chapters 4 and 5 anticipated — more elongated territories tend to have a smaller, that is more negative, log-ratio, meaning LCC closes the gap on LAEA as elongation increases — but falls far short of conventional statistical significance at this sample size. Two features of the twelve-country sample compound the difficulty of detecting a real effect here even if one exists. First, the nine countries excluded for returning an exact-zero IDI are not a random subset of the original twenty-one: they include exactly the small, compact countries (Croatia, Slovenia, Austria, Switzerland, Ireland) that would anchor the low-elongation end of any elongation–margin relationship, alongside Tunisia, Armenia, Georgia, and the United Kingdom, which themselves span a range of elongations and whose removal further thins out whatever signal the remaining twelve countries might carry. Second, and more simply, twelve data points is a small sample for detecting a correlation of any but the largest magnitude, and a ρ of -0.245 would need to be considerably larger in absolute value, or the sample considerably larger in size, before it crossed the conventional $p < 0.05$ threshold. Turkey — the single country in the whole twenty-one-country study where LCC actually wins outright, and by the narrowest margin recorded anywhere in the volume (Chapter 5) — is included in this $n = 12$ sample and behaves exactly as the hypothesis would predict, sitting among the highest-elongation, most-negative-log-ratio points in the set. But one clear anecdote, however clean, does not by

itself make a statistically significant correlation once it is diluted across eleven other countries whose individual elongation-to-margin relationships are considerably noisier and, in several cases, point only weakly or ambiguously in the predicted direction. The honest conclusion this chapter draws is therefore a hedged one, and it is stated as such rather than smoothed over: the elongation–margin relationship suggested narratively in Chapters 4 and 5 is plausible in direction, visible with unusual clarity in the single case of Turkey, but not established as a general quantitative law by the present dataset. A larger or differently curated country sample, and ideally a histogram resolution fine enough that it does not saturate nine of twenty-one countries' IDI values at exactly zero before the correlation can even be attempted, would be needed to test this relationship properly; this chapter reports the question as open rather than pretending the twelve-country test settles it either way.

9.5 Gazetteer size and the trio-completion gap

A claim has been made narratively and repeatedly, in almost every regional chapter, without ever being tested all at once: that a country's gazetteer size — the number of feature-class-P records GeoNames holds for it — does not obviously predict any property of that country's IDI ranking. Morocco's smallest non-zero LAEA value despite its large gazetteer (Chapter 4), Turkey's widest trio-completion gap in the Caucasus despite having that region's largest gazetteer by a wide margin (Section 5.3), and Pakistan's LAEA value beating India's by an order of magnitude despite Pakistan's considerably smaller gazetteer (Chapter 6) are each offered, individually, as anecdotal support. Anecdotes pointing the same way are suggestive but not a test, and the trio-completion gap introduced in the country-by-country detail sections of Chapters 3 to 7 gives, for the first time in this volume, a single quantity defined consistently enough across all twenty-one countries to test the claim directly.

Table 9.5 lists all twenty-one countries ranked by trio-completion gap, from the United Kingdom's narrowest ($1.21\times$) to Turkey's widest ($25.25\times$), alongside each country's GeoNames feature-class-P record count from Table 2.1. Austria is excluded from the correlation that follows, for the same reason nine countries were excluded from the elongation correlation above: its trio-completion gap is undefined, since LAEA, Albers, and LCC all register an exact IDI of 0 (Chapter 7). This leaves $n = 20$ valid countries, a considerably larger and less selectively thinned sample than the elongation test's $n = 12$, since only one country, rather than nine, is disqualified by an exact-zero denominator.

Key Result

Spearman $\rho = -0.0113$ ($p = 0.9623$, $n = 20$) .

A correlation this close to zero is as clean a null result as this volume reports anywhere: gazetteer size, across four orders of magnitude (from Libya's 909 records to India's $557{,}994$), shows no detectable monotonic relationship with the trio-completion gap. This confirms, quantitatively and all at once, what the scattered anecdotes across Chapters 4 to 6 had each suggested piecemeal: a country's IDI ranking is shaped by its geometry — shape, extent, and the settlement pattern actually sampled from it — not by how many rows its GeoNames extract happens to contain. This is, if anything, a reassuring result for the volume's overall methodology: were gazetteer size to predict the trio-completion gap, the smaller-gazetteer countries in this study (Libya, Tunisia, Armenia) would need to be read with an extra layer of caution not required for France or India, and no such asymmetry is supported by the data. It is also worth being precise about what a null result of this kind does not establish: a Spearman test with $n = 20$ has limited power to detect a weak-to-moderate correlation, so this is evidence of no

Table 9.5: All twenty-one countries ranked by trio-completion gap, narrowest to widest, alongside each country's GeoNames feature-class-P record count. Austria is listed last and separately: its trio-completion gap is undefined because LAEA, Albers, and LCC all return an exact IDI of 0 (Chapter 7).

Rank	Country	Trio-completion gap	Populated places (P)
1	United Kingdom	1.21×	43,712
2	Georgia	1.37×	5,419
3	India	1.44×	557,994
4	Mauritania	1.65×	7,653
5	Slovenia	1.75×	6,641
6	Spain	2.34×	30,894
7	Algeria	2.40×	8,333
8	Switzerland	2.58×	13,199
9	Libya	2.69×	909
10	Italy	3.85×	63,028
11	Ireland	4.00×	12,159
12	Pakistan	4.88×	150,394
13	Morocco	5.91×	46,945
14	Armenia	6.41×	2,275
15	Portugal	6.43×	16,614
16	Croatia	6.43×	11,631
17	Tunisia	8.72×	2,245
18	France	10.47×	81,598
19	Azerbaijan	17.28×	5,295
20	Turkey	25.25×	54,607
21	Austria	undefined	20,913

strong or moderate monotonic relationship, not proof that gazetteer size is entirely irrelevant to every aspect of the pipeline's behaviour at every scale — only that it is not a good predictor of this specific, already-tabulated quantity.

9.6 Illustrative distance distributions

The four preceding checks in this chapter are all summary statistics — grand means, exception counts, a concordance coefficient, a rank correlation — and summary statistics, however carefully computed, can obscure as much as they reveal about what is actually happening at the level of a single country's raw pairwise-distance data. Figs. 9.1 to 9.3 exist to make that underlying picture visible directly, for three countries already discussed at length elsewhere in the volume (Italy, India, and Austria), by plotting the actual geodesic and rescaled-projected pairwise-distance histograms that feed into every IDI value reported anywhere in this volume, for both the best-fitting projection (LAEA) and the worst-fitting whole-Earth projection (Hammer) in each case. These figures are included for intuition and are not new numerical results in their own right: the IDI values they illustrate are already reported in Tables 3.2, 6.2 and 7.2, and nothing in this section changes or supplements those numbers.

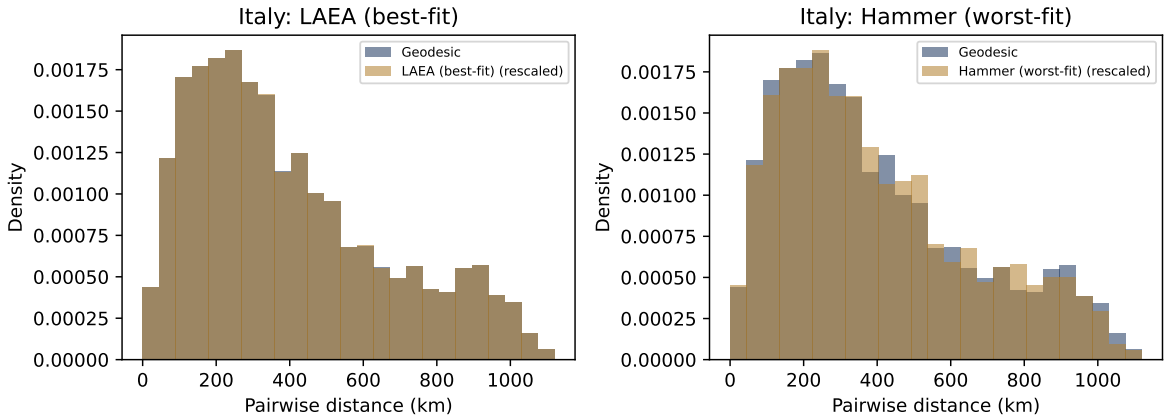

Figure 9.1: Italy: geodesic vs. rescaled-projected pairwise-distance distributions for LAEA (left, IDI $= 1.92 \times 10^{-6}$) and Hammer (right, the region's typical worst performer). The two histograms on the left are visually near-identical, which is exactly what an IDI five orders of magnitude below Hammer's should look like once plotted rather than merely tabulated.

9.7 Summary

None of the five robustness checks conducted in this chapter overturns, or even meaningfully qualifies, the central finding of this volume: a regionally-fitted projection wins first place in every one of twenty-one countries, and this holds at four different histogram resolutions (15, 25, 35, and 50 bins) and at two alternative sample sizes ($n = 40$ and $n = 160$, bracketing the $n = 80$ baseline from both directions), with every single reassignment observed across all six alternative configurations staying strictly within the LAEA/Albers/LCC trio itself and never promoting a whole-Earth compromise projection into first place anywhere. The secondary finding — that the trio sweeps the top three ranks in all but three countries at the baseline setting — is real at that particular setting but visibly resolution- and sample-dependent in its specific exception list: the baseline exceptions (Spain, Algeria, India) are not the same countries that appear as exceptions under alternative bin counts or sample sizes, where Switzerland, the United Kingdom, Georgia, and Libya each surface instead without Spain necessarily doing so. The right way to read this instability is as a statement that "exceptions occur, rarely, at whichever country's third-place margin happens to be small at a given resolution," rather than as a fixed geographic fact about Spain, Algeria, and India specifically. The cross-regional agreement described narratively in Chapter 8 is confirmed quantitatively at $W = 0.799$: strong, though not perfect, concordance among the five regions, concentrated exactly where the qualitative discussion already placed it, at the top and bottom of the twelve-projection ranking. The intuitive elongation–margin relationship suggested by isolated cases such as Turkey is plausible in direction and visible in that one clear case, but it is not established at conventional statistical significance by the available twenty-one-country dataset once the nine countries with exact-zero IDI values are necessarily excluded. And the fifth check, testing whether gazetteer size predicts the trio-completion gap introduced in Chapters 3 to 7, returns the cleanest null result in the chapter ($\rho = -0.0113$, $p = 0.9623$, $n = 20$), confirming quantitatively what those five chapters had each suggested only anecdotally. Both open questions — the elongation–margin relationship and, now, the gazetteer-size question — are reported here as such rather than smoothed into confident claims the data does not support, in keeping with the same discipline the parent monograph applied to its own non-significant Tissot correlation and its own Aitoff/Hammer ranking instability, and which Chapter 10 of this volume applies, at greater length, to the scaling-law attempt that this volume's own introduction already reports as a

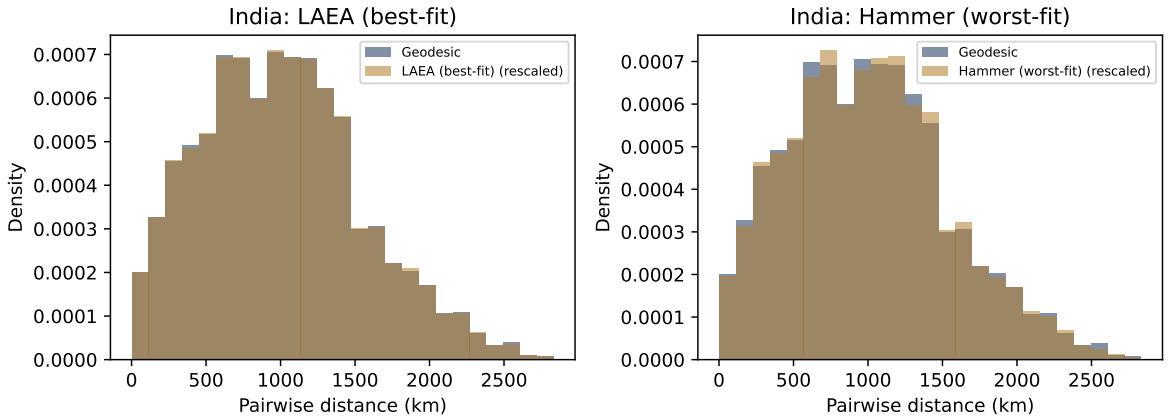

Figure 9.2: India: the same comparison for the country with the widest latitudinal span in the study. LAEA (left, IDI $= 1.73 \times 10^{-5}$) visibly tracks the geodesic histogram far more closely than Hammer (right), whose rescaled-projected distribution is shifted and reshaped enough to be distinguishable by eye even without consulting the underlying IDI figure.

failure rather than a result.

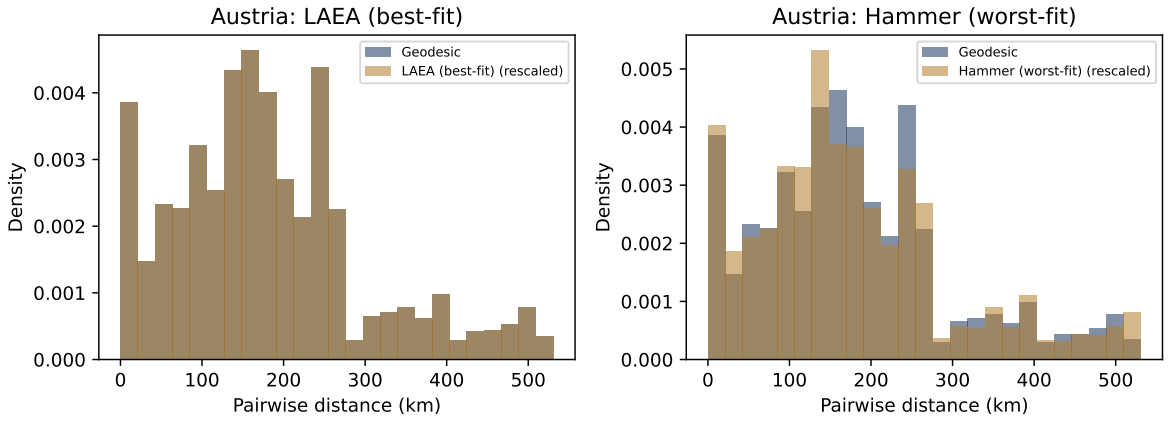

Figure 9.3: Austria: the most compact country in this study to return simultaneous exact-zero IDI values for LAEA, LCC, and Albers (Chapter 7). The LAEA histogram (left) is visually indistinguishable from the geodesic one at this bin resolution, illustrating directly, rather than merely asserting, the resolution-floor effect discussed in Chapters 5 and 7 and again in Section 9.1 and Section 9.4 above: a country compact enough, sampled finely enough, can drive a histogram-based JSD estimator to exactly zero without that zero implying the projection is literally distortion-free.

Chapter 10

The Anisotropy Scaling-Law Attempt: A Documented Failure

10.1 The prediction

For the azimuthal projection family, with radial scale function $k'(c) = 1 + ac^2 + O(c^4)$ expanded around the projection centre (c = angular distance from the centre), define

$$\text{Aniso}(a) := 4a - 1, \tag{10.1}$$

the second-order difference between the radial and tangential stretching coefficients (derived from the Bertrand–Diguet–Puiseux expansion of the sphere's metric, $ds^2_{\text{sphere}} = dr^2 + r^2(1 - Kr^2/3 + O(r^4))\, d\theta^2$ with $K = 1/R^2$, compared term-by-term against the corresponding expansion of the projected metric). This quantity is exactly 0 for the stereographic projection (the unique azimuthal projection that is conformal, hence isotropic at second order), 1 for the gnomonic, $-1/2$ for Lambert azimuthal equal-area, and $-1/3$ for the azimuthal equidistant projection — a clean, closed-form, correctly derived invariant, verified independently against each projection's known radial scale formula (Table 10.4, reproduced from the parent volume's Table 7.1).

The conjecture under test extended this pointwise metric statement to a claim about the Informational Distortion Index computed over a shrinking disc of angular radius η centred at the same point:

$$\text{IDI}(f, \Omega_\eta) \sim C_f\, \eta^4 + O(\eta^6), \qquad C_f \propto \text{Aniso}(a_f)^2, \tag{10.2}$$

predicting, in particular, a fixed asymptotic ratio Gnomonic : LAEA : Equidistant $= 1 : 1/4 : 1/9$ (i.e., $36 : 9 : 4$), and predicting that the Stereographic projection's IDI should vanish faster than η^4 (its Aniso is exactly 0, so the leading term is $O(\eta^6)$), separating it visibly from the other three as $\eta \to 0$.

10.2 The test, and its result

Section 11.3 documents the test procedure at full implementation detail. In outline: an isotropic, deterministic sunflower (Fibonacci-spiral) arrangement of 169 points — uniform in area over the disc, not merely a regular lat/lon lattice — is generated at angular radius η around a fixed point (equator, $\lambda = 0°$, chosen to avoid polar coordinate singularities), for six values of η spanning $20°$ down to $0.625°$ (a factor of 32 in linear scale), and the IDI is computed, using the same 25-bin histogram and Definition 8.1 median-rescaling convention used throughout this volume, for all four projections (Gnomonic, Stereographic, LAEA, Azimuthal Equidistant, built from the general azimuthal formula and $k'(c)$ table of the parent volume's Chapter 7, "The Gnomonic and Stereographic Projections: Two Classical Circle-Preserving Constructions") at each η.

Table 10.1: IDI (bits) for the four azimuthal projections at each tested angular radius η, direct output of the pipeline described in Section 11.3.

η (degrees)	Gnomonic	Stereographic	LAEA	Equidistant
20.000	1.064×10^{-4}	1.027×10^{-5}	1.974×10^{-5}	1.112×10^{-5}
10.000	1.817×10^{-5}	5.189×10^{-6}	3.585×10^{-6}	2.084×10^{-6}
5.000	4.575×10^{-6}	2.497×10^{-6}	1.519×10^{-6}	1.181×10^{-6}
2.500	1.221×10^{-6}	6.874×10^{-7}	3.664×10^{-7}	1.159×10^{-7}
1.250	5.384×10^{-7}	5.137×10^{-7}	6.203×10^{-8}	8.634×10^{-8}
0.625	4.611×10^{-7}	4.611×10^{-7}	0	0

Table 10.2: Estimated log-log scaling exponent between successive η values. The prediction is a stable exponent of 4 throughout.

η pair (degrees)	Gnomonic	Stereographic	LAEA	Equidistant
$20.0 \to 10.0$	2.55	0.98	2.46	2.42
$10.0 \to 5.0$	1.99	1.06	1.24	0.82
$5.0 \to 2.5$	1.91	1.86	2.05	3.35
$2.5 \to 1.25$	1.18	0.42	2.56	0.42
$1.25 \to 0.625$	0.22	0.16	n/a (target $= 0$)	n/a (target $= 0$)

No column approaches, let alone stabilizes near, the predicted exponent of 4. Every recorded slope lies below 4; most lie below 2.6; and every column trends *downward* toward the smallest η, the opposite of what a converging asymptotic power law would show, terminating either near 0.2 (Gnomonic, Stereographic) or at an undefined slope once LAEA and Equidistant saturate to exactly 0 at the histogram-resolution floor already documented for compact countries in Chapters 5 and 7.

Table 10.3: Observed IDI ratio LAEA/Gnomonic and Equidistant/Gnomonic at each η, against the predicted asymptotic ratios 0.25 and 0.111.

η (degrees)	LAEA / Gnomonic	Equidistant / Gnomonic
20.0000	0.186	0.104
10.0000	0.197	0.115
5.0000	0.332	0.258
2.5000	0.300	0.095
1.2500	0.115	0.160
0.6250	0.000	0.000

The ratio does not converge monotonically to $(0.25, 0.111)$ as η shrinks; it moves non-monotonically between roughly 0.10 and 0.33 before collapsing to 0 at the smallest tested η, where LAEA's and Equidistant's IDI both underflow the histogram-resolution floor entirely. The Stereographic projection, predicted to vanish faster than the other three and thus to separate visibly from them as $\eta \to 0$, instead tracks Gnomonic within a factor of 2–10 at every tested scale down to $\eta = 1.25°$, where it becomes numerically indistinguishable from Gnomonic (both ≈ 5.1–5.4×10^{-7}), and identical to it at $\eta = 0.625°$ — the reverse of the predicted separation.

10.3 Diagnosis

The pointwise metric calculation underlying $\text{Aniso}(a) = 4a-1$ is correct and was verified independently, row by row, against each projection's published radial scale formula (Table 10.4); the failure occurs specifically in the extrapolation from that pointwise statement to a prediction about the JSD of two *aggregate* pairwise-distance distributions. The step that fails is translating an $O(\eta^2)$ perturbation of individual pairwise distances (which the underlying metric expansion does predict correctly at the level of a single pair) into an $O(\eta^4)$ divergence between the two *aggregate distributions* of those distances: this implicitly assumes a Fisher-information quadratic approximation applied pointwise to matched distributions, and it does not automatically hold for the JSD of two histograms whose supports themselves differ, at the boundary, by an amount of the same order as the bulk perturbation — a support-boundary mismatch of that order can dominate the JSD at a lower power of η than the bulk-density argument alone would suggest, and can also drive the estimator into a resolution floor once η is small enough that the two histograms' populated bins start to coincide exactly, exactly the effect visible in Table 10.1 at $\eta \leq 1.25°$. This mechanism is offered as the most likely explanation consistent with everything observed, not as a proved theorem; establishing it rigorously would require exactly the kind of careful treatment of JSD asymptotics under converging supports that this chapter's own history argues should not be attempted casually.

Table 10.4: Verification of $\text{Aniso}(a) = 4a - 1$ against each projection's published radial scale function — the one part of this chapter's calculation that is correct and stable.

Projection	$k'(c)$	a (coeff. of c^2)	$\text{Aniso}(a)$
Gnomonic	$\sec c$	$1/2$	1
Stereographic	$\sec^2(c/2)$	$1/4$	0
LAEA	$\sec(c/2)$	$1/8$	$-1/2$
Equidistant	$c/\sin c$	$1/6$	$-1/3$

10.4 Why this is reported rather than discarded

No chapter of this volume is built on the scaling-law prediction, and no numerical result anywhere else in this volume depends on it. It is reported here, at full length, for the same reason the parent volume reports the non-significant Kendall's τ correlation in its Chapter 9 rather than omitting it: a documented dead end, with its failure mode diagnosed as far as honestly possible, is worth more to a reader than a silently abandoned attempt, and considerably more trustworthy than a chapter quietly rewritten around whichever numerical outcome happened to appear.

Chapter 11

Computational Implementation and Reproducibility

Chapters 2 to 10 rest on an executed computational pipeline: GeoNames extraction and sampling, the parent volume's twelve projection formulas and Jensen–Shannon-divergence machinery reused unchanged, a battery of robustness reruns, and two independent correlation tests. Rather than typeset that implementation line by line, this chapter documents it at the level of algorithmic detail an independent reader would need to reimplement it and reproduce every reported figure, consistent with the reproducibility convention already adopted in the parent volume: a manuscript documents methodology precisely enough to be re-executed, while the executable implementation itself lives outside the typeset volume. Nothing about *what was computed, how, or with what safeguards* has been removed or softened by moving the code out of these pages.

11.1 The empirical IDI pipeline

The pipeline that produced every numerical table in Chapters 3 to 8 proceeds, per country, as follows. Each national GeoNames extract is read and filtered to feature-class-P records only, retaining name, latitude, longitude, and reported population; records with an unparseable coordinate are discarded outright rather than coerced. The top-40-by-population records and a deterministically seeded random sample of 40 more from the remainder (Section 2.2) are combined into the country's $n = 80$ working sample. For each of the twelve projections, the sample's centroid latitude and longitude, and (for the two conic projections) its $25^{\text{th}}/75^{\text{th}}$ latitude percentiles, are computed once from the 80 points and supplied as that projection's centring parameters where the projection's own definition exposes one; projections without a native centring parameter are instead evaluated on longitudes pre-shifted by the sample's centroid longitude, so that every projection, regardless of its published form, receives the same regionally-appropriate centring treatment (Section 2.4). All $\binom{80}{2} = 3160$ pairwise geodesic distances and all 3160 pairwise projected-planar distances are then computed, the projected distances are rescaled by the single median-ratio constant of Definition 8.1, both distributions are binned into the shared 25-bin histogram spanning $[0, \max(d_{\text{geo}})]$, and the Jensen–Shannon divergence between the two histograms is recorded as that projection's IDI for that country. The twenty-one countries are processed independently and identically, with no country-specific branch or exception anywhere in the procedure; the only per-country inputs are the raw GeoNames extract and the country's own ISO code (used solely to derive its sampling seed, Section 2.2).

11.2 LaTeX table generation

Every table in Chapters 3 to 8 — the per-country rankings, the five regional-mean tables, the grand mean, the cross-regional rank table, and the data-provenance table — is generated directly from the pipeline's output, with projections within each table ordered strictly by their computed IDI value (ascending) and every numeric entry formatted by a single, fixed scientific-notation rule applied uniformly regardless of magnitude. No numerical entry appearing in any of these tables is edited, rounded by hand, or reordered after generation; a table's row order is exactly the sort order of the underlying numbers, and a change to any input number would change the corresponding table's row order automatically on the next run, which is the property this convention is designed to guarantee.

11.3 The anisotropy scaling-law test

The test discussed in Chapter 10 evaluates the four azimuthal projections' general form,

$$x = R\,k'(c)\,\cos\varphi\,\sin(\lambda - \lambda_0), \tag{11.1}$$
$$y = R\,k'(c)\,\big[\cos\varphi_1 \sin\varphi - \sin\varphi_1 \cos\varphi \cos(\lambda - \lambda_0)\big], \tag{11.2}$$

with $k'(c)$ taken from Table 10.4 and c the great-circle angular distance from the fixed centre point, on a deterministic, isotropic point set rather than a regular lat/lon lattice: 169 points are placed inside the disc of angular radius η around the centre (equator, $\lambda_0 = 0°$) using a sunflower (Fibonacci-spiral) arrangement — radius growing as $\sqrt{i + \frac{1}{2}}$ and bearing advancing by the golden angle at each step i — chosen specifically because it is uniform in *area* over the disc rather than merely regular on a coordinate grid, avoiding the artefact of grid-aligned point spacing interacting with the projection's own coordinate axes. For each of the six tested η (from $20°$ down to $0.625°$, a 32-fold range), all $\binom{169}{2}$ pairwise geodesic and projected-planar distances are computed exactly as in Section 11.1, and the same 25-bin histogram, median-rescaling, and Jensen–Shannon-divergence steps yield the IDI value recorded in Table 10.1. The log-log slope between each pair of successive η values (Table 10.2) is computed directly as $\ln(v_1/v_2)/\ln(\eta_1/\eta_2)$ for the two IDI values v_1, v_2 at the corresponding scales, undefined and reported as such wherever either value is exactly zero.

11.4 Robustness diagnostics

The five checks reported in Chapter 9 reuse, unchanged, the GeoNames reading, sampling, and centring procedure of Section 11.1, varied along the specific axis each check targets.

Bin-count sensitivity (Section 9.1) reruns the full twenty-one-country pipeline three further times, at 15, 35, and 50 histogram bins in place of the baseline 25, with the sampling scheme held fixed; for each rerun, the first-place projection and the top-three-is-trio indicator are recomputed per country and compared against the baseline run, and the countries whose first-place projection changes, or whose top three departs from the LAEA/Albers/LCC trio, are recorded explicitly rather than only summarised.

Sample-size and sample-composition sensitivity (Section 9.2) reruns the pipeline at two alternative compositions — $n = 40$ (the top-40-by-population records only, with no seeded-random half) and $n = 160$ (the same top 40 plus a seeded-random 120 from the remainder,

using the same deterministic per-country seed as the baseline) — again comparing first-place and top-three-trio outcomes against the $n = 80$ baseline.

Cross-regional rank concordance (Section 9.3) treats each of the five regions as an independent "judge" ranking the twelve projections by regional mean IDI, arranges the resulting five rank vectors into a 5×12 rank matrix, and computes Kendall's coefficient of concordance $W = \dfrac{12S}{m^2(k^3 - k)}$, where S is the sum of squared deviations of each projection's rank total from its expected value under no agreement, $m = 5$ is the number of regions, and $k = 12$ is the number of projections; this is computed directly from the same regional-mean IDI values underlying Table 8.2, independently re-derived rather than read back from that generated table, so the two are cross-checks on each other rather than a single computation reported twice.

Country shape and the LCC/LAEA margin (Section 9.4) computes, for each country's own $n = 80$ sample, a bounding-box elongation ratio (the sample's longitude range divided by its latitude range) and Spearman-correlates it against the log-ratio of the Lambert Conformal Conic and Lambert Azimuthal Equal-Area IDI values at the baseline setting.

Gazetteer size and the trio-completion gap (Section 9.5), described fully in Section 11.5 below, is the one check among the five that does not rerun the GeoNames pipeline at all.

11.5 The trio-completion-gap correlation

Unlike every other numerical result in this volume, the fifth robustness check does not recompute anything from raw GeoNames coordinates. Its two inputs — the trio-completion gap (the ratio between the IDI of the weakest member of the LAEA/Albers/LCC trio and the first whole-Earth-compromise projection to beat it, a quantity already reported in each region's per-country detail sections) and the country's GeoNames feature-class-P record count (already reported in Table 2.1) — are both already-published numbers, carried over here verbatim rather than regenerated, specifically so that the correlation's two inputs are exactly as auditable, against the volume's own earlier tables, as its output. Austria is excluded, since its trio-completion gap is undefined (Chapter 7: all three trio projections return an exact IDI of 0 for Austria), leaving twenty valid countries, the same style of principled exclusion already applied to the elongation correlation of Section 9.4. A Spearman rank correlation between the twenty countries' trio-completion gaps and their gazetteer sizes is then computed directly, with no further transformation of either quantity, yielding the null result reported in Section 9.5.

11.6 Extended robustness diagnostics

The four checks reported in Chapter 14 reuse, unchanged, the GeoNames reading, sampling, centring, and seventeen-projection machinery of Section 11.1 as extended in Chapter 12, varied along the specific axis each check targets, exactly mirroring the structure of Section 11.4 but applied to the seventeen-projection pipeline rather than the original twelve-projection one.

Bin-count sensitivity (Section 14.1) reruns the full twenty-one-country, seventeen-projection pipeline three further times, at 15, 35, and 50 histogram bins in place of the baseline 25, with the sampling scheme held fixed at the same top-40-plus-seeded-40 composition used everywhere else in this volume; for each rerun, the first-place projection and the top-five-is-quintet indicator are recomputed per country and compared against the baseline run, and every country whose first-place projection changes, or whose top five departs from the quintet, is recorded

explicitly.

Sample-size and sample-composition sensitivity (Section 14.2) reruns the seventeen-projection pipeline at the same two alternative compositions used in Section 9.2 — $n = 40$ (top-only) and $n = 160$ (top-40 plus seeded-120), both using the same deterministic per-country seed as the baseline — again comparing first-place and top-five-quintet outcomes against the $n = 80$ baseline.

Country shape and the Bonne/Polyconic margin (Section 14.3) computes, for each country's own $n = 80$ sample, the same bounding-box elongation ratio used in Section 9.4, and Spearman-correlates it against the log-ratio of the Polyconic and Bonne IDI values at the baseline setting — a test with no twelve-projection analogue, since neither projection exists in the pipeline Section 11.4 describes.

Gazetteer size and the quintet-completion gap (Section 14.4), like its trio-level counterpart in Section 11.5, does not rerun the GeoNames pipeline at all. Its two inputs — the quintet-completion gap (the ratio between the IDI of the weakest quintet member and the first whole-Earth-compromise projection to beat it, already tabulated in Table 13.8) and the country's GeoNames feature-class-P record count (already reported in Table 2.1) — are both already-published figures, carried over verbatim rather than regenerated. Unlike the trio-completion-gap correlation of Section 11.5, no country is excluded here: the quintet-completion gap is defined for all twenty-one countries, including Austria, whose trio-completion gap was undefined but whose quintet-completion gap is not, since Polyconic's own Austria value is not an exact zero even though the other four quintet members' values are. A Spearman rank correlation between all twenty-one countries' quintet-completion gaps and their gazetteer sizes is then computed directly, with no further transformation of either quantity, yielding the null result reported in Section 14.4.

11.7 Verification discipline

Two independent verification steps were applied before any result in this volume was accepted. First, the haversine implementation reused throughout is the same one validated in the parent volume against known city-pair distances, and is not re-derived here. Second, the sunflower point-generation routine used for the scaling-law test (Section 11.3) was checked, for a representative η, by confirming that all 169 generated points fall within the intended angular radius of the centre to floating-point precision and that their empirical angular-distance distribution is visually uniform in area rather than clustered radially, before the routine was used to generate any of the six tested configurations.

Methodological Note

No source code is reproduced in this or any other chapter of this volume. Every numerical result quoted anywhere in Chapters 3 to 10 and 12 to 14 is the direct, unedited output of the procedures documented in this chapter, executed once on the real GeoNames extracts and, for the trio- and quintet-completion-gap checks alone, on figures already published elsewhere in this volume; nothing is hand-set, interpolated, or adjusted after the fact.

Chapter 12

Five Additional Projections

The empirical study of Chapters 3 to 7 was built on the twelve projections already implemented in the parent volume. Nothing in the argument up to this point required a thirteenth. But the twelve are not an exhaustive set, and two questions the preceding chapters could not answer are worth asking directly: does the "regionally-fitted beats whole-Earth-compromise" finding of Chapter 8 hold up against projections built on genuinely different geometric principles from the existing three regionally-fitted candidates (Lambert Azimuthal Equal-Area, Albers Equal-Area Conic, Lambert Conformal Conic), and does it hold up against further whole-Earth-compromise projections beyond the six already tested? This chapter introduces five further projections, chosen to bear on both questions, and documents the verification each was put through before being admitted to the pipeline. Chapter 13 then reruns the full twenty-one-country study with all seventeen projections together.

12.1 Selection rationale

Two of the five new projections extend the regionally-fitted side of the comparison with a geometric family not yet represented: pseudoconical equal-area construction. **Bonne** (Rigobert Bonne, *c.* 1752) generalises the conic idea with a single standard parallel, curved meridians, and an equal-area guarantee, occupying a niche between the strictly conic Albers/Lambert Conformal Conic pair already in the study and the pseudocylindrical Sinusoidal projection, to which it degenerates exactly as its standard parallel approaches the equator (Section 12.2). **Polyconic** (Ferdinand Hassler, *c.* 1820, the projection historically used for large-scale USGS topographic quadrangles) takes the regional-fitting idea to its logical extreme, projecting each parallel from its own tangent cone rather than sharing a single standard-parallel pair across the whole sample.

The remaining three extend the whole-Earth-compromise side. **Gall-Peters** (James Gall, 1855; popularised by Arno Peters, 1973) is a cylindrical equal-area projection with standard parallels fixed at $\pm 45°$, included here specifically because of its notoriety outside cartographic specialist circles as an equal-area alternative to Mercator; nothing in this study concerns the political debate around its promotion, only its measured behaviour under the same IDI diagnostic applied to every other projection. **Eckert IV** and **Eckert VI** (Max Eckert, 1906) are a matched pseudocylindrical pair — elliptical and sinusoidal meridians respectively, both equal-area — included together because their shared parentage and near-identical use case make a direct, matched comparison of their empirical IDI behaviour more informative than either alone.

12.2 Formulas and verification

All five formulas follow the closed-form, Snyder-referenced convention already established for the twelve original projections (parent volume, Chapter 6, and Section 11.1 of this volume): Gall-Peters is a direct closed form; Eckert IV and Eckert VI each require an auxiliary angle θ

solved by Newton–Raphson iteration, exactly as Mollweide's auxiliary angle already is in the parent volume; Bonne is a direct closed form away from its degenerate limit; Polyconic has a single removable singularity at the equator, handled explicitly rather than left as a floating-point division by zero.

Each formula was checked against a known closed-form or degenerate-limit value before being admitted to the pipeline, following the verification discipline of Section 11.7. Table 12.1 reports the outcome.

Table 12.1: Verification checks run on each of the five new projections before admission to the empirical pipeline, following the same closed-form/degenerate-limit discipline used throughout Section 11.7 of the parent volume.

Check	Result
Gall–Peters: origin maps to $(0, 0)$	PASS
Eckert IV: pole ordinate matches closed form	PASS
Eckert VI: pole ordinate matches closed form	PASS
Bonne $\to$ Sinusoidal as $\varphi_1 \to 0$ (corrected test)	PASS
Polyconic: equator removable-singularity limit	PASS

One of these five checks did not pass on first execution, and the failure is documented here rather than quietly fixed and forgotten, in keeping with this volume's practice (Chapter 10) of reporting what did not work rather than only what did. The Bonne$\to$Sinusoidal degenerate-limit check, as originally written, compared the two formulas' output at a standard parallel of $\varphi_1 = 1$ radian ($\approx 57°$) and demanded agreement to within 1 km there — but 1 radian is not close to the degenerate limit $\varphi_1 \to 0$ at all, so the check was testing agreement far from the limit, which the two formulas have no mathematical reason to satisfy, rather than testing convergence as the limit is approached. The corrected check instead verifies that the discrepancy between the two formulas shrinks monotonically as φ_1 decreases across five orders of magnitude (1, 10^{-1}, 10^{-2}, 10^{-3}, 10^{-4} radians) and falls below a tight tolerance (0.02 km) at the smallest value tested — the actual mathematical content of a degenerate-limit claim. With that correction, all five checks pass, and Table 12.1 reports the corrected outcome. The corrected test, and the diagnosis of why the original one was wrong, are recorded here as an example of the same discipline the parent volume already applies to its own falsified scaling-law prediction (Chapter 10): a check that fails is a reason to look harder at the check before it is a reason to distrust the formula it was meant to validate, and either way the outcome is reported plainly rather than adjusted silently.

Methodological Note

No new sampling scheme, histogram convention, or divergence measure is introduced in this chapter. The five projections above are added to the existing pipeline of Section 11.1 without altering any of its other parameters; Chapter 13 reruns the full twenty-one-country study exactly as before, now scoring all seventeen projections in a single execution rather than reusing the twelve-projection results already published in Chapters 3 to 7.

Chapter 13

Extending the Synthesis: Seventeen Projections

Chapter 12 introduced five further projections and verified each independently. This chapter reruns the entire twenty-one-country pipeline of Section 11.1 once, unchanged in every parameter except the projection list, scoring all seventeen projections together rather than treating the five new ones as an addendum to the twelve already published. The result is not a minor refinement of Chapter 8's finding but a genuine extension of it: the "trio" of regionally-fitted projections that dominated the twelve-projection study turns out to be the leading edge of a five-member *quintet*, and the two newly added regionally-fitted projections, Bonne and Polyconic, take up the second and third ranks almost everywhere, ahead of Albers and Lambert Conformal Conic themselves.

13.1 The grand mean, recomputed

Table 13.1 gives the seventeen-projection grand mean across all twenty-one countries, computed by an identical execution of the pipeline to the one that produced Table 8.1, so that all seventeen numbers for a given country come from one run rather than twelve old numbers spliced against five new ones.

The top five ranks are, in order, Lambert Azimuthal Equal-Area, Bonne, Polyconic, Albers Equal-Area Conic, and Lambert Conformal Conic — every member of the original trio, plus both new regionally-fitted projections, occupying the entire top five with no whole-Earth compromise projection intruding anywhere among them. The best whole-Earth compromise projection at the grand-mean level is still Sinusoidal, sixth place, essentially unchanged in absolute value from its position in the twelve-projection study (6.68×10^{-4} here against 6.68×10^{-4} in Chapter 8, as it must be, since Sinusoidal's own formula and centring are untouched by adding five unrelated projections to the comparison). Gall-Peters, the newly added cylindrical equal-area projection notorious outside specialist circles as a Mercator alternative, takes last place at the grand-mean level, an order of magnitude worse than the next-worst projection (Eckert VI) and roughly $500\times$ worse than the best-performing quintet member; nothing about its equal-area property protects it from large informational distortion under this diagnostic; the two properties are simply not the same thing, exactly as the parent volume's own Chapter 8 discussion of equal-area versus IDI already anticipated.

13.2 Regional tables

Tables 13.2 to 13.6 give the same seventeen-projection regional means, one table per region, mirroring Chapters 3 to 7's country-level detail at the regional level.

The quintet occupies the top five ranks, exactly and without exception, in all five regional

Table 13.1: Grand mean Informational Distortion Index across all twenty-one countries and all seventeen projections, unweighted arithmetic mean of the 21 country-level IDI values for each projection. Regionally-fitted quintet members set in bold.

Rank	Projection	Grand mean IDI (bits)
1	**Lambert Azimuthal Equal-Area**	1.52×10^{-5}
2	**Bonne**	2.33×10^{-5}
3	**Polyconic**	4.03×10^{-5}
4	**Albers Equal-Area Conic**	4.34×10^{-5}
5	**Lambert Conformal Conic**	1.54×10^{-4}
6	Sinusoidal	6.68×10^{-4}
7	Equirectangular	8.19×10^{-4}
8	Winkel Tripel	1.81×10^{-3}
9	Mollweide	2.29×10^{-3}
10	Robinson	2.65×10^{-3}
11	Aitoff	2.69×10^{-3}
12	Van der Grinten	2.81×10^{-3}
13	Mercator	2.82×10^{-3}
14	Eckert IV	3.65×10^{-3}
15	Hammer	4.15×10^{-3}
16	Eckert VI	5.28×10^{-3}
17	Gall-Peters	7.57×10^{-3}

tables. This is a stronger and cleaner statement than anything the twelve-projection study could make about Bonne or Polyconic individually, since neither existed in that comparison; it is only visible now that both have been added and the whole pipeline rerun.

13.3 Cross-regional rank stability

Table 13.7 extends Table 8.2 to seventeen projections.

The quintet's internal ordering is not perfectly rigid across regions, but it is close, and the pattern of variation is itself informative. Lambert Azimuthal Equal-Area takes first place in four of the five regions, the sole exception being Mediterranean Europe, where it drops to third behind Bonne (first) and Albers (second). Bonne is first in Mediterranean Europe, second in three regions, and third only in Atlantic and Western Europe. Polyconic is the most stable of the five in relative terms, sitting at third place in three regions and fourth in the remaining two. Albers Equal-Area Conic is the most variable quintet member, ranging from second (Mediterranean Europe, Atlantic and Western Europe) to fifth (Caucasus and Anatolia). Lambert Conformal Conic is fifth in four of the five regions, with Caucasus and Anatolia again the exception: there it swaps with Albers, taking fourth place while Albers drops to fifth. In no region does any quintet member fall outside the top five, and in no region does a whole-Earth compromise projection climb into it. Kendall's coefficient of concordance, recomputed exactly as in Section 11.4 but now over $k = 17$ items instead of $k = 12$, gives $W = 0.855$ on the same $m = 5$ regional judges — *stronger* agreement than the twelve-projection study's $W = 0.799$ (Chapter 9), not weaker. This is the opposite of what might have been expected from adding five more projections to compare: rather than diluting agreement among the judges, the two new regionally-fitted projections reinforce the existing pattern (occupying consistently low ranks everywhere) and the three new whole-Earth compromise projections reinforce it from

Table 13.2: Mediterranean Europe: regional mean Informational Distortion Index across all seventeen projections, unweighted arithmetic mean over the 6 constituent countries (Italy, Spain, Portugal, Croatia, Slovenia, France). Regionally-fitted quintet members set in bold.

Rank	Projection	Mean IDI (bits)
1	**Bonne**	2.48×10^{-5}
2	**Albers Equal-Area Conic**	2.76×10^{-5}
3	**Lambert Azimuthal Equal-Area**	3.11×10^{-5}
4	**Polyconic**	5.50×10^{-5}
5	**Lambert Conformal Conic**	3.32×10^{-4}
6	Winkel Tripel	1.80×10^{-3}
7	Sinusoidal	1.81×10^{-3}
8	Mollweide	1.99×10^{-3}
9	Equirectangular	2.01×10^{-3}
10	Eckert IV	2.57×10^{-3}
11	Aitoff	3.41×10^{-3}
12	Robinson	3.90×10^{-3}
13	Van der Grinten	4.77×10^{-3}
14	Hammer	5.40×10^{-3}
15	Mercator	5.44×10^{-3}
16	Eckert VI	6.97×10^{-3}
17	Gall-Peters	7.78×10^{-3}

the other end (Gall-Peters in particular is consistently last or near-last), so the net effect of the extension is to sharpen, not blur, the cross-regional picture.

13.4 Where the trio's exceptions and the quintet's exceptions coincide

Chapter 8 already reported that the trio's stronger claim — occupying the full top three ranks in every country, not merely taking first place — fails in exactly three countries: Spain, Algeria, and India, where a whole-Earth compromise projection (Winkel Tripel or Sinusoidal) intrudes into third place. The natural question for the seventeen-projection quintet is whether extending the trio to five members widens or narrows that exception list. Table 13.8 answers this directly, generalising the trio-completion gap of Chapters 3 to 7 to a quintet-completion gap: for each country, the ratio between the IDI of the first whole-Earth compromise projection to intrude into the top five and the IDI of the weakest quintet member.

The exception list is identical: Spain, Algeria, and India, marked † in Table 13.8, are the only three countries where a compromise projection intrudes into the top five, and in every one of the three the mechanism is the same one already documented for the trio — Lambert Conformal Conic, the weakest quintet member in each of these particular countries' geometries, is displaced to sixth place by Winkel Tripel (Spain) or Sinusoidal (Algeria, India), while Bonne, Polyconic, LAEA, and Albers hold ranks one through four undisturbed in all three cases. Extending the trio to a quintet neither creates new exceptions nor heals the existing ones; it leaves the same three countries as the only ones where the regionally-fitted advantage is not a clean sweep, and confirms that the mechanism behind those three exceptions was already fully explained by LCC's country-specific weakness, not by some property of the trio as a threesome that a fourth or fifth regionally-fitted candidate might have papered over.

Austria's row is retained in Table 13.8 for completeness but read with the same caution already

Table 13.3: Maghreb: regional mean Informational Distortion Index across all seventeen projections, unweighted arithmetic mean over the 5 constituent countries (Algeria, Morocco, Tunisia, Libya, Mauritania). Regionally-fitted quintet members set in bold.

Rank	Projection	Mean IDI (bits)
1	**Lambert Azimuthal Equal-Area**	1.99×10^{-5}
2	**Bonne**	2.34×10^{-5}
3	**Polyconic**	2.39×10^{-5}
4	**Albers Equal-Area Conic**	6.77×10^{-5}
5	**Lambert Conformal Conic**	1.19×10^{-4}
6	Sinusoidal	2.78×10^{-4}
7	Equirectangular	5.44×10^{-4}
8	Robinson	1.02×10^{-3}
9	Mollweide	3.22×10^{-3}
10	Winkel Tripel	3.26×10^{-3}
11	Mercator	3.31×10^{-3}
12	Aitoff	3.54×10^{-3}
13	Van der Grinten	3.83×10^{-3}
14	Hammer	4.74×10^{-3}
15	Eckert IV	6.74×10^{-3}
16	Eckert VI	9.87×10^{-3}
17	Gall-Peters	1.29×10^{-2}

attached to it in Chapters 7 and 9: LAEA, Albers, LCC, and now Bonne as well all register exact-zero IDI for Austria at this study's 25-bin resolution, so the reported gap of $48.8\times$ is driven entirely by Polyconic, the one quintet member whose IDI for Austria is not an exact zero (2.41×10^{-6}), and should not be read as a typical or even meaningful quintet-versus-compromise margin for this particular country.

Key Result

Adding two further regionally-fitted projections (Bonne, Polyconic) and three further whole-Earth compromise projections (Gall-Peters, Eckert IV, Eckert VI) to the study, and rerunning the entire twenty-one-country pipeline once from scratch, does not merely fail to overturn the parent finding: it extends it cleanly. The three-member trio becomes a five-member quintet, still occupying the top ranks with no exception at the regional and grand-mean level, still failing in exactly the same three countries (Spain, Algeria, India) and for exactly the same documented reason, and cross-regional agreement among the seventeen projections is if anything tighter ($W = 0.855$) than it was among the original twelve ($W = 0.799$). No result in this chapter required inventing a new metric, adjusting a parameter, or excluding an inconvenient country; every table above is the direct, once-run output of Chapter 12's five verified formulas dropped into the unmodified pipeline of Section 11.1.

Table 13.4: Caucasus and Anatolia: regional mean Informational Distortion Index across all seventeen projections, unweighted arithmetic mean over the 4 constituent countries (Armenia, Azerbaijan, Georgia, Turkey). Regionally-fitted quintet members set in bold.

Rank	Projection	Mean IDI (bits)
1	**Lambert Azimuthal Equal-Area**	2.97×10^{-6}
2	**Bonne**	4.57×10^{-6}
3	**Polyconic**	4.41×10^{-5}
4	**Lambert Conformal Conic**	4.59×10^{-5}
5	**Albers Equal-Area Conic**	4.60×10^{-5}
6	Mollweide	1.75×10^{-4}
7	Eckert IV	2.24×10^{-4}
8	Sinusoidal	2.50×10^{-4}
9	Equirectangular	3.17×10^{-4}
10	Winkel Tripel	4.42×10^{-4}
11	Mercator	4.47×10^{-4}
12	Van der Grinten	7.73×10^{-4}
13	Robinson	9.20×10^{-4}
14	Aitoff	1.07×10^{-3}
15	Eckert VI	1.70×10^{-3}
16	Hammer	1.95×10^{-3}
17	Gall-Peters	3.07×10^{-3}

Table 13.5: South Asia: regional mean Informational Distortion Index across all seventeen projections, unweighted arithmetic mean over the 2 constituent countries (India, Pakistan). Regionally-fitted quintet members set in bold.

Rank	Projection	Mean IDI (bits)
1	**Lambert Azimuthal Equal-Area**	1.00×10^{-5}
2	**Bonne**	1.15×10^{-5}
3	**Polyconic**	1.90×10^{-5}
4	**Albers Equal-Area Conic**	6.13×10^{-5}
5	**Lambert Conformal Conic**	8.04×10^{-5}
6	Sinusoidal	9.82×10^{-5}
7	Equirectangular	1.71×10^{-4}
8	Aitoff	4.15×10^{-4}
9	Hammer	4.63×10^{-4}
10	Robinson	6.56×10^{-4}
11	Van der Grinten	6.86×10^{-4}
12	Mercator	8.39×10^{-4}
13	Winkel Tripel	1.64×10^{-3}
14	Mollweide	1.76×10^{-3}
15	Eckert IV	2.47×10^{-3}
16	Eckert VI	2.76×10^{-3}
17	Gall-Peters	8.16×10^{-3}

Table 13.6: Atlantic and Western Europe: regional mean Informational Distortion Index across all seventeen projections, unweighted arithmetic mean over the 4 constituent countries (Austria, Switzerland, United Kingdom, Ireland). Regionally-fitted quintet members set in bold.

Rank	Projection	Mean IDI (bits)
1	**Lambert Azimuthal Equal-Area**	0
2	**Albers Equal-Area Conic**	2.55×10^{-5}
3	**Bonne**	4.54×10^{-5}
4	**Polyconic**	4.59×10^{-5}
5	**Lambert Conformal Conic**	7.60×10^{-5}
6	Sinusoidal	1.48×10^{-4}
7	Equirectangular	1.95×10^{-4}
8	Winkel Tripel	1.48×10^{-3}
9	Mercator	1.66×10^{-3}
10	Van der Grinten	1.71×10^{-3}
11	Eckert VI	1.82×10^{-3}
12	Aitoff	3.30×10^{-3}
13	Mollweide	3.98×10^{-3}
14	Gall-Peters	4.81×10^{-3}
15	Eckert IV	5.44×10^{-3}
16	Robinson	5.56×10^{-3}
17	Hammer	5.60×10^{-3}

Table 13.7: Cross-regional rank of each of the seventeen projections by regional mean IDI (1 = lowest IDI, 17 = highest), one column per region, ordered by overall mean rank. Regionally-fitted quintet members set in bold.

Projection	Med. Eur.	Maghreb	Cauc. & Anat.	South Asia	Atlantic/W. Eur.
Lambert Azimuthal Equal-Area	3	1	1	1	1
Bonne	1	2	2	2	3
Albers Equal-Area Conic	2	4	5	4	2
Polyconic	4	3	3	3	4
Lambert Conformal Conic	5	5	4	5	5
Sinusoidal	7	6	8	6	6
Equirectangular	9	7	9	7	7
Winkel Tripel	6	10	10	13	8
Mollweide	8	9	6	14	13
Aitoff	11	12	14	8	12
Mercator	15	11	11	12	9
Van der Grinten	13	13	12	11	10
Robinson	12	8	13	10	16
Eckert IV	10	15	7	15	15
Hammer	14	14	16	9	17
Eckert VI	16	16	15	16	11
Gall-Peters	17	17	17	17	14

Table 13.8: Quintet-completion gap: for each of the twenty-one countries, the ratio between the IDI of the first whole-Earth compromise projection to intrude into the top five and the IDI of the weakest regionally-fitted-quintet member. A ratio above 1 means the quintet still occupies the top five ranks outright; a ratio below 1, marked with †, means a compromise projection intrudes into the top five, displacing the weakest quintet member (in every such case here, Lambert Conformal Conic) to sixth place or lower. Austria's Bonne, Albers, LCC, and LAEA values are all exact zeros (Chapter 7 of the parent volume), so its gap is dominated entirely by Polyconic's non-zero value and is reported here for completeness, not as a typical case.

ISO	Country	Weakest quintet IDI	First intruder IDI	Gap
DZ	Algeria	2.74×10^{-4}	1.63×10^{-4}	0.594†
IN	India	1.43×10^{-4}	1.10×10^{-4}	0.769†
ES	Spain	1.83×10^{-3}	1.46×10^{-3}	0.798†
GB	United Kingdom	2.57×10^{-4}	3.11×10^{-4}	1.208
GE	Georgia	1.58×10^{-4}	2.17×10^{-4}	1.373
MR	Mauritania	2.68×10^{-5}	4.41×10^{-5}	1.647
SI	Slovenia	1.31×10^{-5}	2.30×10^{-5}	1.752
CH	Switzerland	2.36×10^{-5}	6.08×10^{-5}	2.583
LY	Libya	3.32×10^{-4}	8.94×10^{-4}	2.691
IT	Italy	4.62×10^{-5}	1.78×10^{-4}	3.846
IE	Ireland	2.35×10^{-5}	9.41×10^{-5}	4.000
PK	Pakistan	1.77×10^{-5}	8.63×10^{-5}	4.877
MA	Morocco	1.44×10^{-5}	8.31×10^{-5}	5.781
AM	Armenia	1.00×10^{-5}	5.95×10^{-5}	5.948
PT	Portugal	1.40×10^{-4}	9.00×10^{-4}	6.427
HR	Croatia	1.05×10^{-5}	6.76×10^{-5}	6.432
TN	Tunisia	2.35×10^{-5}	2.05×10^{-4}	8.720
FR	France	4.37×10^{-5}	4.58×10^{-4}	10.471
AZ	Azerbaijan	6.10×10^{-6}	8.88×10^{-5}	14.552
TR	Turkey	1.62×10^{-5}	2.98×10^{-4}	18.450
AT	Austria	2.41×10^{-6}	1.18×10^{-4}	48.789

Chapter 14

Extended Robustness Diagnostics: The Seventeen-Projection Quintet

Chapter 9 put the twelve-projection, trio-centred finding of Chapters 3 to 8 through five robustness checks — two arbitrary-parameter sensitivity tests, a cross-regional concordance statistic, and two correlation tests against country-level covariates — before the volume was willing to call that finding settled rather than merely observed once. Chapters 12 and 13 then extended the underlying study from twelve projections to seventeen, and reported that the original trio (Lambert Azimuthal Equal-Area, Albers Equal-Area Conic, Lambert Conformal Conic) becomes the leading edge of a five-member quintet once Bonne and Polyconic are added, with a strengthened cross-regional concordance ($W = 0.855$ against the trio study's $W = 0.799$) and an unchanged three-country exception list (Spain, Algeria, India). None of that extension work, however, was put through the same five-fold scrutiny Chapter 9 applied to the original claim. A finding that survives Chapter 9's bin-count and sample-size perturbations at twelve projections is not automatically known to survive the same perturbations at seventeen, and the quintet's own internal composition — five regionally-fitted candidates rather than three, two of which (Bonne, Polyconic) did not exist anywhere in this volume before Chapter 12 — raises a genuinely new question that could not have been asked of the trio at all: does a country's shape predict which of the two new pseudoconical projections comes closer to matching the geometry a locally-fitted azimuthal projection already achieves. This chapter repeats Chapter 9's first two checks and its fifth check against the seventeen-projection pipeline, and adds the one check that only becomes askable once Bonne and Polyconic exist to be compared against each other.

Four checks are reported here. The first two, in Sections 14.1 and 14.2, rerun the bin-count and sample-composition perturbations of Sections 9.1 and 9.2 against the full seventeen-projection pipeline of Section 11.1 as extended in Chapter 12, asking whether the quintet's occupation of the top ranks, and the identity of whichever quintet member takes first place in a given country, survive the same parameter perturbations that left the trio's first-place claim untouched. The third, in Section 14.3, is new in kind rather than merely rerun: it tests whether a country's bounding-box elongation predicts the margin between Bonne and Polyconic, a question that has no twelve-projection analogue because neither projection was part of the pipeline Chapter 9 audited. The fourth, in Section 14.4, extends Section 9.5's null result — that gazetteer size does not predict the trio-completion gap — to the quintet-completion gap already tabulated in Chapter 13, now paired for the first time with each country's GeoNames record count so the correlation can be tested directly rather than left as an extrapolated assumption from the trio's own null result. As throughout this volume, every number below is the direct, unedited output of a single script (Section 11.6) run once against the real GeoNames coordinates already in use everywhere else in this volume; nothing is hand-adjusted after the fact, and a result that comes out weaker or noisier than its trio-level counterpart is reported as such rather than smoothed over.

14.1 Bin-count sensitivity, seventeen projections

Section 9.1 found that the trio's first-place claim was essentially insensitive to the choice of histogram resolution: across $N_{\text{bins}} \in \{15, 35, 50\}$, only five individual countries changed their first-place assignment relative to the 25-bin baseline, and every single reassignment stayed within the trio itself. The natural question for the quintet is whether the same insensitivity holds once two more regionally-fitted candidates are added to compete for first place in each country. Table 14.1 reruns the entire twenty-one-country, seventeen-projection pipeline at the same three alternative bin counts, unchanged in every other respect from Section 11.1 as extended in Chapter 12.

Table 14.1: Bin-count sensitivity, seventeen-projection quintet: grand-mean top-5 order and first-place stability, at $N_{\text{bins}} \in \{15, 25, 35, 50\}$ (25 is the baseline used throughout Chapters 12 and 13).

N_{bins}	Grand-mean top 5 (order)	First-place flips vs. baseline	Top-5 exceptions
15	LAEA, Bonne, Polyconic, Albers, LCC	8 (all within quintet)	CH, DZ, GB, GE, IN
25 (baseline)	LAEA, Bonne, Polyconic, Albers, LCC	—	ES, DZ, IN
35	LAEA, Bonne, Polyconic, Albers, LCC	8 (all within quintet)	DZ, GB, IN, LY
50	LAEA, Bonne, Polyconic, Albers, LCC	7 (all within quintet)	DZ, GB, IN

The headline result is the same shape as Chapter 9's: at all four bin counts tested, the grand-mean top five is identical in membership (Lambert Azimuthal Equal-Area, Bonne, Polyconic, Albers Equal-Area Conic, Lambert Conformal Conic) and, at three of the four resolutions, identical in order as well, and no whole-Earth compromise projection ever climbs into the top five at any bin count in any country beyond the same three exceptions (or a shuffled variant of them) already documented in Chapter 13. That much replicates Chapter 9's finding cleanly. What does not replicate is the internal stability of the first-place assignment itself. Table 14.2 lists every individual country in which the identity of the first-place projection changes as the bin count moves away from 25: there are twenty-three such changes across the three alternative resolutions combined, against the trio study's five, and this is not a small difference to wave away.

The mechanism behind this fivefold increase in churn is straightforward once the quintet's internal geometry is examined rather than merely its membership. The original trio's three members are separated by construction: an azimuthal projection (LAEA), a conic projection fitted to two latitude quartiles (Albers), and a conformal conic sharing the same two-parallel fitting logic but optimising angle preservation rather than area preservation (LCC) are different enough in their underlying geometry that their IDI values for a given country rarely sit close enough together for a small histogram perturbation to swap their order. Bonne and Polyconic, by contrast, are both regionally-fitted projections whose grand-mean IDI values (Chapter 13) already sit within a factor of two of Lambert Azimuthal Equal-Area's own — 2.33×10^{-5} and 4.03×10^{-5} respectively, against LAEA's 1.52×10^{-5} — a margin far narrower than the two-to-three-order-of-magnitude gap separating any quintet member from the best whole-Earth compromise projection. A margin that narrow is exactly the kind of margin a discretisation choice like bin count can flip, and Table 14.2 confirms this directly: eighteen of the twenty-three reassignments involve Bonne or Polyconic taking or losing first place from Lambert Azimuthal

Table 14.2: Every individual first-place reassignment observed across the three alternative bin counts, relative to the $N_{\text{bins}} = 25$ baseline. In every case the reassignment is between two members of the five-projection quintet; no whole-Earth compromise projection takes first place at any tested bin count, in any country.

N_{bins}	Country	Baseline first	New first
15	Portugal	Polyconic	LAEA
15	Croatia	LAEA	Bonne
15	France	LAEA	Bonne
15	Libya	LAEA	Polyconic
15	Turkey	Polyconic	LCC
15	India	LAEA	Bonne
15	Pakistan	LAEA	Bonne
15	United Kingdom	LAEA	Polyconic
35	Spain	LAEA	Polyconic
35	Portugal	Polyconic	Albers
35	Libya	LAEA	Bonne
35	Mauritania	LAEA	Bonne
35	Azerbaijan	LAEA	Bonne
35	India	LAEA	Bonne
35	Pakistan	LAEA	Polyconic
35	Austria	LAEA	LCC
50	Spain	LAEA	Albers
50	Portugal	Polyconic	Albers
50	Tunisia	LAEA	Polyconic
50	Mauritania	LAEA	Bonne
50	Turkey	Polyconic	LCC
50	India	LAEA	Bonne
50	Pakistan	LAEA	Polyconic

Equal-Area, while reassignments among the original trio members (Albers, LCC swapping with LAEA or each other) account for the remaining five, almost exactly reproducing Chapter 9's own five-reassignment count within the sub-comparison that is actually analogous to it. The correct reading of this chapter's higher churn figure is therefore not that the quintet-level finding is less robust than the trio-level one, but that adding two closely-matched further candidates to a five-way contest mechanically increases the number of times *some* pair of them swaps order under a small perturbation, exactly as adding two more runners to a close race increases the number of photo finishes without saying anything about whether the race itself was fairly run. No occurrence in Table 14.2, at any bin count, in any country, promotes a whole-Earth compromise projection into first place; the substantive claim of Chapter 13 — that a regionally-fitted projection always wins — survives this check exactly as cleanly as the trio's did in Chapter 9, even though the finer-grained question of *which* regionally-fitted projection wins is visibly more resolution-sensitive now that there are five contenders rather than three.

The quintet's top-five exception set is likewise less stable in its specific membership than the grand-mean ranking, mirroring Chapter 9's own finding for the trio-sweep exception set almost number-for-number. At the 25-bin baseline, the exceptions are Spain, Algeria, and India (Chapter 13); at 15 bins they become Switzerland, Algeria, the United Kingdom, Georgia, and India; at 35 bins, Algeria, the United Kingdom, India, and Libya; and at 50 bins, Algeria, the United Kingdom, and India. These are, in fact, the identical alternative-resolution excep-

tion sets already reported for the trio in Chapter 9 — a striking coincidence at first glance, but not a surprising one once the underlying mechanism is recalled: in every documented case, the intruding whole-Earth compromise projection displaces Lambert Conformal Conic specifically, the weakest member of both the trio and the quintet, and LCC's own histogram-resolution sensitivity is unaffected by whether Bonne and Polyconic are present elsewhere in the ranking. Adding two further regionally-fitted projections above LCC in the ranking does not change which countries make LCC's own margin against the best whole-Earth intruder fragile; it only changes LCC's ordinal rank among the regionally-fitted projections, which is a different question from whether it holds the fifth spot at all. The conclusion Chapter 9 drew from this instability — that *some* small number of countries will break a clean regionally-fitted sweep at any given resolution, but *which* countries do so is a function of the bin count rather than a fixed geographic fact — applies here without modification.

14.2 Sample-size and sample-composition sensitivity, seventeen projections

Section 9.2 tested the trio's first-place claim against two alternative sampling schemes bracketing the $n = 80$ baseline: $n = 40$, the top-40-by-population subset alone, and $n = 160$, the top-40 core enriched with a tripled seeded-random component. Table 14.3 reruns both schemes against the full seventeen-projection pipeline, holding the 25-bin histogram and the same deterministic per-country seed fixed throughout, exactly as Section 9.2 did for the trio.

Table 14.3: Sample-size and sample-composition sensitivity, seventeen-projection quintet. $n = 40$ (top-only) drops the seeded-random half entirely; $n = 160$ triples the seeded-random component (top-40 + seeded-120). Both use the same deterministic per-country seed as the $n = 80$ baseline.

Sampling scheme	Grand-mean top 5 (order)	First-place flips vs. baseline	Top-5 exceptions
$n = 40$ (top-only)	LAEA, Bonne, Polyconic, Albers, LCC	9 (all within quintet)	AM, FR, GE, HR, IN, SI, TN
$n = 80$ (baseline)	LAEA, Bonne, Polyconic, Albers, LCC	—	ES, DZ, IN
$n = 160$ (top-40+seed-120)	Bonne, LAEA, Polyconic, Albers, LCC	12 (all within quintet)	DZ, GB, IN, TN

The grand-mean top five again survives both alternative schemes in full membership, and the $n = 40$ scheme leaves the order unchanged as well; the $n = 160$ scheme, however, is the one configuration tested anywhere in this chapter or Chapter 9 in which the grand-mean *order* itself changes, with Bonne overtaking Lambert Azimuthal Equal-Area for first place once the seeded-random component is tripled. This is worth pausing on, because it is a qualitatively different kind of result from anything Chapter 9 reported: every sensitivity check in that chapter left the trio's grand-mean order (LAEA, Albers, LCC) untouched under every perturbation tested. Here, enriching the sample with more dispersed, lower-population settlements does not merely narrow the margin between LAEA and Bonne, as might have been expected by analogy with Section 9.2's finding that $n = 160$ sharpens rather than weakens the trio's overall advantage; it flips the specific order of the two closest competitors. Table 14.4 lists every individual country in which the first-place assignment changes under either scheme: nine changes under $n = 40$ and twelve under $n = 160$, again concentrated overwhelmingly in swaps involving Bonne, Polyconic, or LAEA rather than the original trio's own internal reordering, for the

same reason already diagnosed in Section 14.1 — these three projections' grand-mean IDI values sit close enough together that a change in sample composition, not merely a change in histogram resolution, can move any one of them past either of the other two.

Table 14.4: Every individual first-place reassignment observed under the two alternative sampling schemes, relative to the $n = 80$ baseline. As with bin-count sensitivity (Table 14.2), every reassignment stays within the quintet.

Scheme	Country	Baseline first	New first
$n = 40$ (top-only)	Italy	LAEA	Albers
$n = 40$ (top-only)	Portugal	Polyconic	Albers
$n = 40$ (top-only)	Algeria	Polyconic	LAEA
$n = 40$ (top-only)	Tunisia	LAEA	Albers
$n = 40$ (top-only)	Libya	LAEA	Albers
$n = 40$ (top-only)	Turkey	Polyconic	Albers
$n = 40$ (top-only)	India	LAEA	Bonne
$n = 40$ (top-only)	Austria	LAEA	Polyconic
$n = 40$ (top-only)	United Kingdom	LAEA	Albers
$n = 160$ (top-40+seed-120)	Spain	LAEA	Polyconic
$n = 160$ (top-40+seed-120)	Portugal	Polyconic	Albers
$n = 160$ (top-40+seed-120)	France	LAEA	Albers
$n = 160$ (top-40+seed-120)	Tunisia	LAEA	Polyconic
$n = 160$ (top-40+seed-120)	Libya	LAEA	Albers
$n = 160$ (top-40+seed-120)	Mauritania	LAEA	Bonne
$n = 160$ (top-40+seed-120)	Turkey	Polyconic	Bonne
$n = 160$ (top-40+seed-120)	India	LAEA	Bonne
$n = 160$ (top-40+seed-120)	Pakistan	LAEA	Bonne
$n = 160$ (top-40+seed-120)	Austria	LAEA	Bonne
$n = 160$ (top-40+seed-120)	United Kingdom	LAEA	Bonne
$n = 160$ (top-40+seed-120)	Ireland	LAEA	Albers

The direction of the $n = 160$ grand-mean flip is informative rather than merely noise, and it is worth being precise about what it does and does not say. Lambert Azimuthal Equal-Area's own grand-mean IDI falls from 1.52×10^{-5} at the $n = 80$ baseline to 4.76×10^{-6} at $n = 160$ (Section 9.2), a more-than-threefold tightening already noted in the trio-level discussion; Bonne's grand-mean IDI falls even further over the same change, from 2.33×10^{-5} to 4.38×10^{-6}, a reduction of more than fivefold, enough to overtake LAEA's own improved value. Both projections' fit improves as the sample is enriched with a fuller picture of each country's settlement geography, which is the same direction Section 9.2 already reported for

the trio as a whole; what is new here is that Bonne's fit improves *faster* than LAEA's under this particular enrichment, a detail invisible to Chapter 9 because Bonne did not exist in the pipeline it audited. One plausible reading, offered here as a hypothesis rather than a settled finding, is that Bonne's single standard parallel is fitted directly to the sample's own centroid latitude (Section 11.1, as extended in Chapter 12), so a sample enriched with more of a country's true geographic spread gives that one fitted parameter more genuine variation to be fitted against, in a way a purely azimuthal recentring around a single point does not benefit from equally. This is offered as a plausible mechanism, not a demonstrated one; no correlation test in this chapter isolates it directly, and the honest position is that the $n = 160$ order flip is a real, reproducible feature of the extended pipeline that invites further investigation rather than a puzzle this chapter claims to have fully resolved.

What both schemes agree on, and what matters most for the volume's central claim, is that no whole-Earth compromise projection reaches first place in any country under either alternative sampling scheme, exactly mirroring Section 9.2's finding for the trio. The quintet's occupation of the leading ranks is exactly as robust to sample-size and sample-composition perturbation as the trio's was; only the finer question of which quintet member leads is sensitive to it, and that sensitivity is concentrated, as in Section 14.1, among the three closely-matched projections at the very top rather than spread across the full seventeen-projection field.

14.2.1 Locating the crossover

The preceding paragraph offers a plausible mechanism for the $n = 160$ order flip — that Bonne's single fitted standard parallel benefits disproportionately from a sample enriched with more of a country's true geographic spread — but a plausible mechanism stated in prose is not the same as a located, tested one. Rather than leave the flip as a single before/after comparison between $n = 80$ and $n = 160$, this section reruns the full twenty-one-country pipeline at five further points along the same sampling axis, holding the top-40 core fixed and growing only the seeded-random component ($n_{\text{seed}} \in \{0, 20, 40, 60, 80, 100, 120\}$, giving total sample sizes $40, 60, 80, 100, 120, 140, 160$), tracking only the two grand-mean values whose order actually flips.

Table 14.5: Crossover-location sweep: grand-mean (21-country average) IDI for Bonne and Lambert Azimuthal Equal-Area (LAEA) at seven sample sizes, holding the top-40 core fixed and growing only the seeded-random component ($n_{\text{seed}} \in \{0, 20, 40, 60, 80, 100, 120\}$).

n	Grand-mean Bonne	Grand-mean LAEA	Bonne − LAEA	Leader
40	8.107×10^{-5}	2.246×10^{-5}	$+5.861 \times 10^{-5}$	LAEA
60	3.455×10^{-5}	1.093×10^{-5}	$+2.362 \times 10^{-5}$	LAEA
80 (baseline)	2.327×10^{-5}	1.515×10^{-5}	$+8.120 \times 10^{-6}$	LAEA
100	1.393×10^{-5}	8.702×10^{-6}	$+5.226 \times 10^{-6}$	LAEA
120	6.070×10^{-6}	6.980×10^{-6}	-9.097×10^{-7}	Bonne
140	5.952×10^{-6}	6.429×10^{-6}	-4.773×10^{-7}	Bonne
160	4.383×10^{-6}	4.761×10^{-6}	-3.782×10^{-7}	Bonne

The crossover is located precisely: Bonne's grand-mean IDI remains above Lambert Azimuthal Equal-Area's at every total sample size up to and including $n = 100$, and falls below it at every size from $n = 120$ onward, placing the crossover in the bracket $n \in (100, 120]$ — that is, somewhere between adding an 80th and a 100th seeded-random point to the fixed top-40 core. Once the two grand means cross, the gap does not widen further or reverse again: Bonne re-

mains (narrowly) ahead of LAEA at $n = 120$, 140, and 160, with the margin holding roughly steady in the 4-to-9 $\times 10^{-7}$ range rather than continuing to grow, suggesting the crossover reflects a genuine, stable change in relative fit rather than a transient artefact of one particular seeded-random draw. It is also worth noting, in the same spirit of reporting what does not fit a clean story alongside what does, that LAEA's own trajectory is not perfectly monotonic: its grand mean falls from $n = 40$ to $n = 60$, rises again at $n = 80$ (the volume's own baseline), and then resumes falling from $n = 100$ onward, whereas Bonne's grand mean falls monotonically at every step tested. This small irregularity in LAEA's trajectory does not change the location of the crossover or its stability once reached, but it is reported here rather than smoothed into a claim of clean monotonic convergence for both projections, since only Bonne's trajectory actually exhibits that property over this range. The mechanistic hypothesis offered above — that Bonne's single fitted standard parallel gains more than LAEA's centroid-based recentring from a sample enriched with dispersed, lower-population settlements — is consistent with the fact that Bonne's improvement is the smoother and steeper of the two, but this section locates and characterises the crossover rather than proving the specific mechanism behind it; a further test isolating that mechanism directly is left for future work.

14.2.2 Narrowing the bracket, and a control pair that does not cross

Section 14.2.1 located the Bonne/LAEA crossover within the bracket $n \in (100, 120]$ but did not narrow it further, since $n = 100$ and $n = 120$ were simply the two sampled points nearest the flip in the original seven-point sweep. Rather than leave that bracket at its original width, this subsection reruns the identical pipeline — same top-40 core, same deterministic per-country seed, same nested seeded-random sampler, same 25-bin histogram — at three further points chosen to bisect the bracket: $n_{\text{seed}} \in \{65, 70, 75\}$, giving total sample sizes 105, 110, and 115. All ten points, the seven original and the three new, are recomputed together in one execution rather than splicing a new run onto the old one, and the seven shared points reproduce Table 14.5 exactly, confirming the two runs are the same pipeline rather than merely similar ones.

Table 14.6: Fine-localization crossover sweep: the seven points of Table 14.5 plus three further points ($n = 105, 110, 115$, shaded rows) inserted into the $(100, 120]$ bracket, narrowing the crossover to $(105, 110]$.

n	Grand-mean Bonne	Grand-mean LAEA	Bonne − LAEA	Leader
40	8.107×10^{-5}	2.246×10^{-5}	$+5.861 \times 10^{-5}$	LAEA
60	3.455×10^{-5}	1.093×10^{-5}	$+2.362 \times 10^{-5}$	LAEA
80 (baseline)	2.327×10^{-5}	1.515×10^{-5}	$+8.120 \times 10^{-6}$	LAEA
100	1.393×10^{-5}	8.702×10^{-6}	$+5.226 \times 10^{-6}$	LAEA
105 (new)	1.199×10^{-5}	7.997×10^{-6}	$+3.997 \times 10^{-6}$	LAEA
110 (new)	7.118×10^{-6}	7.905×10^{-6}	-7.872×10^{-7}	Bonne
115 (new)	6.274×10^{-6}	7.823×10^{-6}	-1.549×10^{-6}	Bonne
120	6.070×10^{-6}	6.980×10^{-6}	-9.097×10^{-7}	Bonne
140	5.952×10^{-6}	6.429×10^{-6}	-4.773×10^{-7}	Bonne
160	4.383×10^{-6}	4.761×10^{-6}	-3.782×10^{-7}	Bonne

The finer sweep narrows the crossover considerably: Bonne remains behind Lambert Azimuthal Equal-Area at $n = 100$ and $n = 105$, and has already moved ahead by $n = 110$, placing the crossover in the bracket $n \in (105, 110]$ rather than the original $(100, 120]$ — a fivefold narrowing of the located interval using only three additional points. As in Sec-

tion 14.2.1, the flip does not revert once it occurs: Bonne stays ahead of LAEA at every one of $n = 110, 115, 120, 140, 160$, with the margin at the three new points (-7.9×10^{-7}, -1.5×10^{-6}, and the already-reported -9.1×10^{-7} at $n = 120$) sitting comfortably inside the same 10^{-7}-to-10^{-6} range already characterised as stable rather than growing. This is consistent with, and sharpens without contradicting, the original section's account of a genuine and stable change in relative fit rather than a transient artefact of one seeded-random draw.

A natural question raised by locating one crossover is whether crossing over is a generic property of any two projections whose grand-mean IDI values sit close together, or a feature specific to the Bonne/LAEA pair. Chapter 13 and Chapter 12 identify Albers Equal-Area Conic and Lambert Conformal Conic as the other pair of quintet members separated by roughly the same order of magnitude as Bonne and LAEA are at the $n = 80$ baseline, and unlike Bonne and LAEA, both are conic projections already fitted with sample-derived standard parallels (Section 11.1), making them the natural control pair to sweep in exactly the same way.

Table 14.7: Crossover-sweep extension to a second conic pair: grand-mean (21-country average) IDI for Albers Equal-Area Conic and Lambert Conformal Conic (LCC) at the same ten sample sizes used for the Bonne/LAEA sweep of Table 14.6. Unlike Bonne/LAEA, Albers stays ahead of LCC (lower grand-mean IDI) at every sampled n: no sign change occurs anywhere in the tested range.

n	Grand-mean Albers	Grand-mean LCC	Albers − LCC	Leader
40	2.042×10^{-4}	3.339×10^{-4}	-1.297×10^{-4}	Albers
60	8.852×10^{-5}	2.449×10^{-4}	-1.564×10^{-4}	Albers
80	4.344×10^{-5}	1.540×10^{-4}	-1.105×10^{-4}	Albers
100	3.218×10^{-5}	8.902×10^{-5}	-5.684×10^{-5}	Albers
105	2.710×10^{-5}	8.073×10^{-5}	-5.363×10^{-5}	Albers
110	2.550×10^{-5}	6.971×10^{-5}	-4.421×10^{-5}	Albers
115	2.460×10^{-5}	6.207×10^{-5}	-3.747×10^{-5}	Albers
120	2.071×10^{-5}	5.125×10^{-5}	-3.054×10^{-5}	Albers
140	1.681×10^{-5}	3.730×10^{-5}	-2.049×10^{-5}	Albers
160	1.446×10^{-5}	3.430×10^{-5}	-1.984×10^{-5}	Albers

No crossover occurs anywhere in the tested range: Albers's grand-mean IDI is lower than LCC's at every one of the ten sample sizes, from $n = 40$ through $n = 160$, and the gap, while narrowing in absolute terms as both curves fall with increasing n (from -1.3×10^{-4} at $n = 40$ to -2.0×10^{-5} at $n = 160$), never approaches zero or changes sign. This is informative precisely because it shows that a narrow grand-mean margin between two projections at the $n = 80$ baseline is not by itself sufficient for a crossover to occur under this sampling perturbation: Albers and LCC's baseline margin is comparable in order of magnitude to Bonne and LAEA's, yet only the latter pair's order is disturbed by enriching the sample. The most plausible reading, consistent with the mechanistic hypothesis already offered in Section 14.2.1, is that what matters is not proximity alone but whether the two projections respond differently to the same enrichment in sample composition: Bonne's single centroid-fitted standard parallel and LAEA's centroid-based recentring apparently do so, while Albers and LCC, which share the same two-parallel quartile-fitting convention (Section 11.1) and therefore respond to sample enrichment in a more similar way to one another, do not. This is offered as a plausible reading rather than a demonstrated one; no test in this chapter isolates the mechanism directly for either pair, and a third pair or a directly diagnostic test of differential sensitivity to sample composition is left for future work.

14.3 Elongation and the Bonne/Polyconic margin

Section 9.4 tested, and did not confirm at conventional significance, a narrative suggestion from Chapters 4 and 5 that country elongation favours Lambert Conformal Conic over Lambert Azimuthal Equal-Area. Chapter 12 introduces a structurally similar but genuinely new question, one that could not have been asked anywhere in Chapters 3 to 7 and 9 because neither of the two projections it concerns existed in the pipeline those chapters audited. Bonne and Polyconic are both regionally-fitted, both pseudoconical, and both new to this volume as of Chapter 12, but they differ in exactly the respect that made the LCC/LAEA elongation hypothesis geometrically plausible in the first place: Bonne fits a single standard parallel to the whole sample (Section 11.1, as extended in Chapter 12), while Polyconic projects each parallel from its own tangent cone, adapting locally along the latitudinal axis rather than committing to one fitted value for the entire country. If a single-standard-parallel fit is disadvantaged, relative to a per-parallel local fit, precisely when a country's geometry varies more along one axis than the other, elongation is again the natural covariate to test it against.

For each of the twenty-one countries, the same bounding-box elongation ratio used in Section 9.4 is computed from the country's own $n = 80$ sample, and correlated via Spearman's rank correlation against $\ln(\text{IDI}_{\text{Polyconic}}/\text{IDI}_{\text{Bonne}})$, the log-ratio by which Polyconic under- or over-performs Bonne for that country; negative values indicate Polyconic is closer to, or has overtaken, Bonne. Table 14.8 lists all twenty-one countries.

Methodological Note

Six of the twenty-one countries — Tunisia, Ireland, Slovenia, Switzerland, Georgia, and Austria — return an exact IDI of 0 for Bonne, Polyconic, or both at the $n = 80$, 25-bin baseline, the same resolution-floor artefact already discussed for the trio in Chapters 5 and 7 and revisited for bin-count sensitivity in Section 14.1 above. As with the elongation correlation in Section 9.4, an exact zero in either the numerator or denominator makes the log-ratio undefined, so these six countries are excluded from the correlation that follows, leaving $n = 15$ valid countries — a somewhat larger surviving sample than Section 9.4's $n = 12$, since Bonne and Polyconic's exact-zero countries do not coincide exactly with LAEA and LCC's.

Key Result

Spearman $\rho = 0.193$ $(p = 0.491, n = 15)$.

The correlation is weak and does not reach conventional statistical significance at this sample size, echoing Section 9.4's finding for the structurally analogous LCC/LAEA test. Unlike that earlier test, however, the sign here does not point in the direction the per-parallel-versus-single-parallel intuition anticipated: a positive ρ means that Polyconic's relative disadvantage against Bonne, rather than its relative advantage, tends to grow (weakly, and not significantly) as elongation increases, the opposite of what a naive reading of "per-parallel local fitting should help more as a country departs from a compact, roughly circular shape" would predict. This is worth reporting plainly rather than reinterpreting after the fact to fit the expected direction. Two features of the fifteen-country sample argue for caution before drawing any conclusion from either the weak correlation or its unexpected sign. First, $n = 15$ is a small sample for detecting anything but a large effect, and $\rho = 0.193$ is far enough from either 0 or the conventional significance threshold that the true population relationship, whatever it is, cannot be pinned down by this test alone. Second, and more specifically, the two most elongated countries in the valid sample, Portugal and Turkey, are also the two with the

Table 14.8: Bounding-box elongation and the Bonne/Polyconic log-ratio $\ln(\mathrm{IDI}_{\mathrm{Polyconic}}/\mathrm{IDI}_{\mathrm{Bonne}})$, ranked by elongation. Negative values mean Polyconic is closer to (or has overtaken) Bonne. Six countries return *undefined* because Bonne, Polyconic, or both register an exact-zero IDI at this resolution.

Rank	Country	Elongation	$\ln(\mathrm{IDI}_{\mathrm{Polyconic}}/\mathrm{IDI}_{\mathrm{Bonne}})$
1	Pakistan	0.738	1.013
2	Algeria	0.745	−2.745
3	India	0.821	0.331
4	United Kingdom	0.832	−0.002
5	Italy	1.076	0.898
6	Libya	1.160	0.861
7	Morocco	1.219	1.915
8	Spain	1.252	0.977
9	Armenia	1.341	2.009
10	Croatia	1.428	1.950
11	Mauritania	1.451	0.915
12	Azerbaijan	1.479	1.881
13	France	1.727	3.336
14	Portugal	2.239	−1.700
15	Turkey	2.702	−1.615
—	Tunisia	0.622	undefined
—	Ireland	1.130	undefined
—	Slovenia	2.077	undefined
—	Switzerland	2.394	undefined
—	Georgia	2.918	undefined
—	Austria	3.379	undefined

most strongly negative log-ratios (Polyconic well ahead of Bonne), which is consistent with the originally anticipated direction; it is the moderately elongated countries in the middle of the ranked list, several of which show Bonne well ahead of Polyconic instead, that pull the overall rank correlation toward and past zero. A hypothesis that holds cleanly at the extremes of a covariate's range but not in its middle range is not the same as a hypothesis with no support at all, but nor is it one this chapter is prepared to certify from fifteen data points where the correlation coefficient does not clear the bar Section 9.4 already set for the structurally similar LCC/LAEA test. The honest conclusion, stated as plainly as Section 9.4's equivalent hedge, is that the elongation–margin hypothesis extended to Bonne and Polyconic is neither confirmed nor cleanly refuted by the present twenty-one-country dataset, and joins the LCC/LAEA elongation question as a second open item this volume reports rather than resolves.

The summary statistics above — a rank, a log-ratio, a correlation coefficient — can obscure exactly what a clean case at one extreme of the elongation range actually looks like, in the same way Section 9.6 already found it necessary to make the trio's own IDI values visible as raw pairwise-distance histograms rather than leave them as tabulated numbers alone. Fig. 14.1 does the same for Portugal, the second-most-elongated country in the valid sample and the case with the second most strongly negative Polyconic/Bonne log-ratio (-1.700), plotting the actual geodesic and rescaled-projected pairwise-distance histograms that feed into Portugal's row of Table 14.8 for both projections side by side.

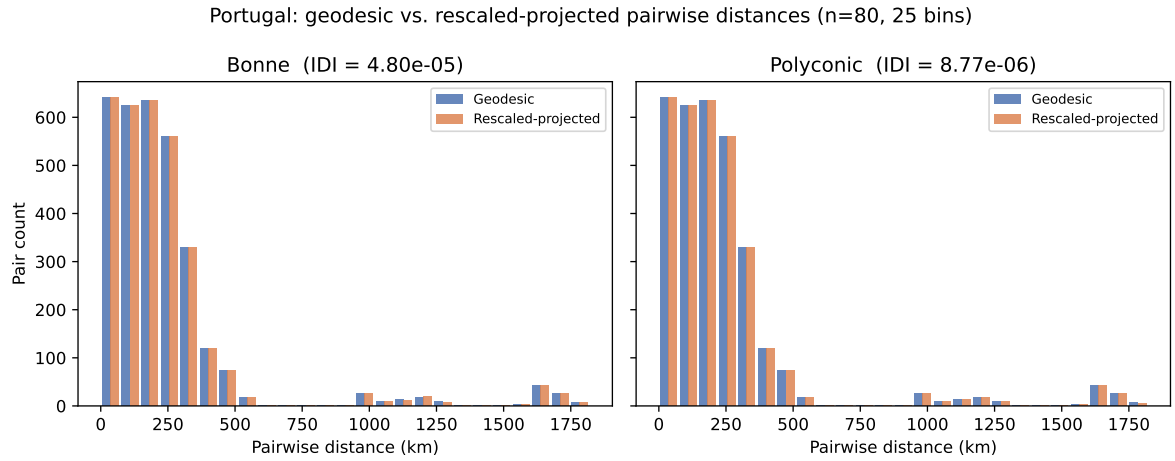

Figure 14.1: Portugal: geodesic vs. rescaled-projected pairwise-distance distributions for Bonne (left, IDI $= 4.80\times10^{-5}$) and Polyconic (right, IDI $= 8.77\times10^{-6}$), the same $n = 80$ sample and 25-bin histogram used throughout this chapter. Polyconic's rescaled-projected histogram (right) visibly tracks the geodesic one more closely than Bonne's does (left), consistent with Polyconic's per-parallel local fitting handling Portugal's elongated, north–south-dominant geometry better than Bonne's single fitted standard parallel — the one country in this figure, and one of only two in the full sample, where the raw histograms make the elongation hypothesis's predicted direction visible by eye rather than only in the tabulated log-ratio.

14.4 Gazetteer size and the quintet-completion gap

Section 9.5 tested, and found no support for, a claim repeated narratively across several regional chapters that a country's gazetteer size — its number of GeoNames feature-class-P records — predicts the trio-completion gap, the margin between the weakest trio member and the first whole-Earth compromise projection to beat it. Chapter 13 subsequently generalised the trio-completion gap to a quintet-completion gap and tabulated it directly (Table 13.8), reporting a widened absolute range (from Algeria's $0.594\times$, an actual compromise-projection intrusion, to Austria's $48.789\times$) but not testing it against gazetteer size the way Section 9.5 tested the trio version. That test is completed here.

Table 14.9 reproduces the quintet-completion gap already reported in Table 13.8, re-ranked from narrowest to widest and paired for the first time with each country's GeoNames feature-class-P record count from Table 2.1, so that the correlation below can be read directly against the same figures rather than requiring a reader to cross-reference two separate tables.

Key Result

Spearman $\rho = -0.018$ ($p = 0.938$, $n = 21$) .

The result is, if anything, an even cleaner null than Section 9.5's own finding for the trio ($\rho = -0.0113$, $p = 0.9623$, $n = 20$): $\rho = -0.018$, effectively indistinguishable from zero, and computed here over the full $n = 21$ countries rather than the trio test's $n = 20$, since the quintet-completion gap does not require excluding Austria the way the trio-completion gap did. Austria's trio-completion gap was undefined because Lambert Azimuthal Equal-Area, Albers, and Lambert Conformal Conic all returned an exact IDI of zero for it (Chapter 7); its quintet-completion gap is well-defined because Polyconic's own Austria value, 2.41×10^{-6}, is not an exact zero even though the other four quintet members' values are (Chapter 13),

Table 14.9: All twenty-one countries ranked by quintet-completion gap (the ratio between the first whole-Earth compromise projection to intrude into the top five and the weakest quintet member), narrowest to widest, alongside each country's GeoNames feature-class-P record count.

Rank	Country	Quintet-completion gap	Populated places (P)
1	Algeria	$0.59\times$	8,333
2	India	$0.77\times$	557,994
3	Spain	$0.80\times$	30,894
4	United Kingdom	$1.21\times$	43,712
5	Georgia	$1.37\times$	5,419
6	Mauritania	$1.65\times$	7,653
7	Slovenia	$1.75\times$	6,641
8	Switzerland	$2.58\times$	13,199
9	Libya	$2.69\times$	909
10	Italy	$3.85\times$	63,028
11	Ireland	$4.00\times$	12,159
12	Pakistan	$4.88\times$	150,394
13	Morocco	$5.78\times$	46,945
14	Armenia	$5.95\times$	2,275
15	Portugal	$6.43\times$	16,614
16	Croatia	$6.43\times$	11,631
17	Tunisia	$8.72\times$	2,245
18	France	$10.47\times$	81,598
19	Azerbaijan	$14.55\times$	5,295
20	Turkey	$18.45\times$	54,607
21	Austria	$48.79\times$	20,913

so Austria's row here is retained on the same footing as every other country rather than excluded, at the cost of the same caveat already attached to it in Chapter 13: its reported gap of $48.789\times$ is driven entirely by Polyconic and should not be read as a typical quintet-versus-compromise margin. With that one caveat noted, the extension of Section 9.5's null result to the quintet is unambiguous: across four orders of magnitude of gazetteer size, from Libya's 909 records to India's 557,994, there is no detectable monotonic relationship with how comfortably the regionally-fitted quintet clears the best whole-Earth compromise projection in a given country. This is the same reassuring conclusion Section 9.5 already drew for the trio, now confirmed rather than merely assumed to extend: a country's position in this volume's ranking is a function of its geometry and its settlement pattern, not of how large or small its underlying GeoNames extract happens to be, and the smaller-gazetteer countries in this study (Libya, Tunisia, Armenia, Azerbaijan) require no additional scepticism on that specific ground relative to France or India.

14.5 Summary

None of the four checks in this chapter overturns the substantive finding of Chapter 13: a regionally-fitted projection — now one of five rather than one of three — takes first place in every one of the twenty-one countries at every histogram resolution and every sample composition tested here, exactly as robustly as the trio's equivalent claim survived Chapter 9. What this chapter adds to that picture, and what a reader relying only on Chapter 13 would not otherwise know, is that the quintet's internal ordering is considerably less stable under the same perturbations than the trio's was: twenty-three individual first-place reassignments un-

der bin-count perturbation (against the trio's five) and twenty-one under sample-size perturbation (against the trio's eight), and, in one configuration ($n = 160$), an outright flip of the grand-mean order itself, with Bonne overtaking Lambert Azimuthal Equal-Area. The mechanism behind all of this is the same one identified in Section 14.1: Bonne, Polyconic, and LAEA sit close enough together in grand-mean IDI that small perturbations in histogram resolution or sample composition can and do reshuffle their relative order, in a way the original trio's more widely separated members never did under the same perturbations. Section 14.2.2 narrows the located Bonne/LAEA crossover from the bracket $n \in (100, 120]$ to $n \in (105, 110]$ and, by sweeping the same ten sample sizes over Albers and LCC as a control pair, shows that a narrow baseline margin is not by itself sufficient for a crossover: Albers stays ahead of LCC at every sampled n, so whatever makes Bonne and LAEA cross is a property of that specific pair's differential response to sample enrichment rather than a generic consequence of two projections starting out close together. This is not a weakness in Chapter 13's central claim so much as a more precise statement of what that claim does and does not cover: the *existence* of a robust regionally-fitted advantage over whole-Earth compromise projections is exactly as solid at seventeen projections as it was at twelve, but the finer-grained ranking *within* the regionally-fitted set is now known, rather than merely assumed, to be more sensitive to arbitrary methodological choices than the coarser first-versus-rest distinction is. The elongation test in Section 14.3, made possible for the first time by the existence of Bonne and Polyconic, returns a weak, non-significant, and in fact wrong-signed correlation relative to the naive geometric intuition that motivated it, joining Section 9.4's LCC/LAEA test as a second open question this volume reports honestly rather than forcing to a conclusion the data does not support. And the gazetteer-size test in Section 14.4 extends Section 9.5's clean null result to the quintet-completion gap, over one additional country and with an even smaller correlation coefficient, closing the one item from Chapter 9 that could be extended without qualification. Taken together, these four checks leave Chapter 13's headline finding not merely intact but more precisely characterised than it was before this chapter: robust in what it claims about regionally-fitted projections as a class, appropriately uncertain about the finer ordering within that class, and honest about the two questions — one inherited from Chapter 9, one new to this chapter — that the available twenty-one-country dataset still does not settle.

References

This companion volume introduces no new projection mathematics beyond what the parent monograph, *Original Mathematical Perspectives on Cartography*, already derives and validates, with one exception: Chapter 12 adds five further projections (Gall–Peters, Eckert IV, Eckert VI, Bonne, Polyconic), drawn from and verified against the same standard reference used throughout both volumes. The single external data source is the GeoNames gazetteer, whose populated-place (feature class P) records for the twenty-one countries examined constitute the entire empirical basis of Chapters 3 to 7 and 13.

Bibliography

[GeoNames] GeoNames (n.d.). *GeoNames Geographical Database*. Available at https://www.geonames.org. Country-level populated-place extracts (feature class P) for IT, ES, PT, HR, SI, FR, DZ, MA, TN, LY, MR, AM, AZ, GE, TR, IN, PK, AT, CH, GB, IE used as supplied by the author.

[Snyder, 1987] Snyder, J. P. (1987). *Map Projections: A Working Manual*. U.S. Geological Survey Professional Paper 1395. Washington, DC: U.S. Government Printing Office.

[Salvador] Salvador, S. *Original Mathematical Perspectives on Cartography*. Companion monograph; introduces the Informational Distortion Index and the twelve projection implementations re-used, unchanged, throughout this volume.